ONTAP hardware systems documentation

Install and maintain

NetApp
August 29, 2025

This PDF was generated from https://docs.netapp.com/us-en/ontap-systems/index.html on August 29, 2025. Always check docs.netapp.com for the latest.

Table of Contents

ONTAP hardware systems documentation

Release notes

What's new for ONTAP hardware systems

Learn what's new for ONTAP hardware systems. For additional support information, see ONTAP hardware and software compatibility and hardware no longer supported.

May 2025

ASA C30

The new ASA C30 system extends high-performance, intelligent, and comprehensive data management capabilities to more customers and workloads.

Learn more about the ASA r2 systems.

April 2025

FAS50

The FAS50 system offers the lowest cost per Gigabyte for secondary storage and secure cyber vault, plus increased performance for secondary workloads with faster backups, higher IOPS, and low latency of 5–10ms.

Learn more about FAS systems.

February 2025

ASA A20, ASA A30, and ASA A50

The entry-level ASA A20 and midrange ASA A30 and A50 systems make block storage available to companies of every size for mission-critical apps like databases and virtual machines.

Learn more about the ASA r2 systems.

December 2024

AFF A20, AFF A30, and AFF A50

The new AFF A20, A30, and A50 hardware systems extend high-performance, intelligent, and comprehensive data management capabilities to more customers and workloads.

The systems offer real-time ML-based ransomware detection, seamless cloud integration, and unmatched performance for mission-critical workloads including AI, VMware, databases, and analytics with support for block, file, and object storage.

Learn more about AFF A-Series systems.

AFF C30, AFF C60, and AFF C80

The new AFF C30, AFF C60, and AFF C80 hardware systems make the performance and efficiency gains of flash more accessible to businesses by providing an industry-leading 1.5PB of storage capacity in two-rack deployments.

The solution offers exceptional density with 60TB drives, increased performance, and improved I/O flexibility.

Learn more about AFF C-Series systems.

NS224 drive shelf with NSM100B shelf modules

Enhance your data center capabilities with the NS224 and the new NSM100B NVMe storage expansion module. Designed as a direct replacement for the NSM100 module, the new NSM100B shelf module integrates seamlessly into your existing setup. It supports both direct-attached and switch-attached configurations of NS224 shelves, offering exceptional flexibility to optimize your storage system's efficiency and scalability.

September 2024

ASA r2 systems

The new ASA r2 hardware systems (ASA A1K, ASA A70, and ASA A90) deliver a unified hardware and software solution that creates a simplified experience specific to the needs of SAN-only customers.

Learn more about the ASA r2 systems.

FAS70 and FAS90

The new FAS70 and FAS90 hardware systems deliver affordable, yet high-performing backup storage, enabling a secure cyber vault for recovery from ransomware attacks.

Learn more about FAS systems.

ONTAP hardware and software compatibility

ONTAP storage systems are compatible with several versions of ONTAP software. Learn about the ONTAP versions that your storage systems and drive shelves support.

Full configuration support and limitations for ONTAP hardware systems are available in NetApp Hardware Universe. Details about known issues, limitations, and upgrade cautions in ONTAP 9 software are available in the ONTAP 9 Release Notes.

AFF systems

The AFF A-Series and AFF C-Series systems provide a robust, scale-out platform tailored for virtualized environments. They can be deployed as standalone systems or as high-performance tiers in NetApp ONTAP configurations. Learn more about AFF A-Series systems and AFF C-Series systems.

AFF A-Series and AFF C-Series systems are supported beginning in the following ONTAP releases.

ONTAP 9.16.1	• AFF A20 • AFF A30 • AFF A50 • AFF C30 • AFF C60 • AFF C80
ONTAP 9.15.1	• AFF A1K • AFF A70 • AFF A90
ONTAP 9.12.1P1	• AFF A150 • AFF C250 • AFF C400 • AFF C800
ONTAP 9.10.1	• AFF A900
ONTAP 9.8	• AFF A250
ONTAP 9.7	• AFF A400 • AFF A800

ASA r2 systems

ASA r2 systems deliver a unified hardware and software solution that creates a simplified experience specific to the needs of SAN-only customers. Learn more about the ASA r2 systems.

ASA r2 systems are supported beginning in the following ONTAP releases.

ONTAP 9.16.1	• ASA A20 • ASA A30 • ASA A50 • ASA C30

ONTAP 9.16.0	• ASA A1K • ASA A70 • ASA A90

ASA systems

ASA A-Series and ASA C-Series systems deliver a simplified and dedicated SAN experience that provides continuous data availability for enterprise mission critical databases and other SAN workloads using FCP or iSCSI protocol. Learn more about the ASA systems.

ASA systems are supported beginning in the following ONTAP releases.

ONTAP 9.13.1P1	• ASA C250 • ASA C400 • ASA C800
ONTAP 9.13.1	• ASA A150 • ASA A250 • ASA A400 • ASA A900
ONTAP 9.8	• ASA AFF A250 • ASA AFF A800
ONTAP 9.7	• ASA AFF A400

FAS systems

FAS systems deliver efficient and secure secondary storage - the ultimate solution for tiering, backup, and disaster recovery. Learn more about FAS systems.

FAS systems are supported beginning in the following ONTAP releases.

ONTAP 9.16.1	• FAS50
ONTAP 9.15.1	• FAS70 • FAS90
ONTAP 9.13.1	• FAS2820
ONTAP 9.11.1	• FAS9500
ONTAP 9.10.1P3	• FAS9500

ONTAP 9.7	• FAS2750 • FAS8300 • FAS8700

Drive shelves

Drive shelves are specifically designed for NetApp AFF, ASA, and FAS systems and help deliver the performance, resiliency, and flexibility your digital transformation needs.

Drive shelves are available beginning with the following ONTAP releases.

ONTAP 9.16.1	• DCM3 for SAS-3 shelves • NS224 with NSM100B modules
ONTAP 9.6	NS224 shelf with NSM100 modules

Hardware no longer supported

The following systems, shelves, and switches are no longer supported as of the specified version of ONTAP.

For current unsupported hardware, see NetApp Hardware Universe .

ONTAP hardware systems

System	Support discontinued from…
• AFF A300 • FAS8200	ONTAP 9.17.1
• AFF A320 • AFF A700s	ONTAP 9.15.1
• AFF A200 • FAS2650 • FAS2620	ONTAP 9.12.1
• AFF AFF8020, AFF8040, AFF8060, AFF8080 • FAS8020, FAS8040, FAS8060, FAS8080 • FAS2520, FAS2552, FAS2554	ONTAP 9.9.1

Shelves

Shelf module	Support discontinued from…
IOM6 6Gb/s SAS module, used in: • DS2246 • DS4246 • DS4486	ONTAP 9.16.1

Switches

Switch	Support discontinued from…
NetApp CN1610 switch	ONTAP 9.13.1
• Cisco 5596UP • Cisco 5596T	ONTAP 9.11.1

Related information

- Supported Cisco Ethernet Switches
- Supported NetApp Ethernet Switches
- End of Availability Platforms

Quick start for ONTAP hardware systems

To get up and running with ONTAP hardware systems, you install hardware components, cable your hardware, and configure your storage in ONTAP.

If your system is in a MetroCluster configuration, go to the MetroCluster doc site and follow the install instructions applicable to your MetroCluster configuration type.

Use the following workflow to deploy your storage system when it is not set up in a MetroCluster configuration.

Install switches

Install your switches in the rack or cabinet. Access the following instructions for your switch model.

Cluster switches

- Install BES-53248 switch
- Install Cisco Nexus 9336C-FX2 switch
- Install NVIDIA SN2100 switch

Storage switches

- Install Cisco Nexus 9336C-FX2 switch
- Install NVIDIA SN2100 switch

Shared switches

- Install Cisco Nexus 9336C-FX2 switch

Install the storage system and storage shelves

Install your storage system and storage shelves in the cabinet or rack. Access the install and setup instructions for your platform model.

- AFF systems
- ASA systems
- ASA r2 systems
- FAS systems

Connect cables

Cable the controllers to your network and then cable the controllers to your shelves. The install and setup instructions for your platform model include instructions for cabling the controller ports to your network and to your switches.

Set up your ONTAP cluster

After you have installed and set up your controllers and switches, you must complete configuring your storage in ONTAP. Access the following instructions according to your deployment configuration.

- For ONTAP deployments, see Configure ONTAP.
- For ONTAP with MetroCluster deployments, see Configure Metrocluster with ONTAP.

AFF systems

AFF A-Series systems

AFF A1K systems

Install and setup

Installation and configuration workflow - AFF A1K

To install and configure your AFF A1K system, you review the hardware requirements, prepare your site, install and cable the hardware components, power on the system, and set up your ONTAP cluster.

Review installation requirements

Review the equipment and tools needed to install your storage system and storage shelves and review the lifting and safety precautions.

Prepare to install the AFF A1K storage system

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

Install the hardware for the AFF A1K storage system

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system in the cabinet or telco rack. Next, slide the shelves onto the rails. Finally, attach cable management devices to the rear of the storage system for organized cable routing.

Cable the controllers and storage shelves for the AFF A1K storage system

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

Power on the AFF A1K storage system

Before you power on the controllers, power on each NS224 shelf and assign a unique shelf ID to ensure each shelf is uniquely identified within the setup, connect the laptop or console to the controller, and then connect the controllers to the power sources.

Set up your cluster

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

Review the equipment needed and the lifting precautions for your AFF A1K storage system and storage shelves.

Equipment needed for install

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

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

Lifting precautions

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

Storage system weight

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

An A1K storage system can weigh up to 62.83 lbs (28.5 kg). To lift the storage system, use two people or a hydraulic lift.

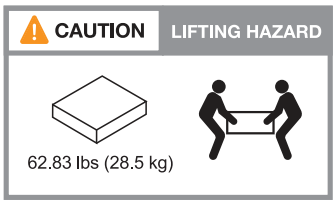

Shelf weight

Take the necessary precautions when moving or lifting your shelf.

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

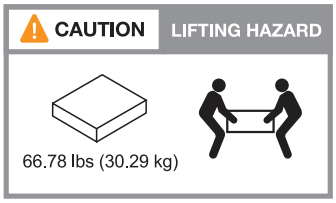

Related information

- Safety information and regulatory notices
  :a1k!:

What's next?

After you've reviewed the hardware requirements, you prepare to install your AFF A1K storage system.

Prepare to install your AFF A1K storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the system to access support benefits.

Step 1: Prepare the site

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

Steps

1. Use NetApp Hardware Universe to confirm that your site meets the environmental and electrical requirements for your storage system.
2. Make sure you have adequate cabinet or rack space for your storage system, shelves, and any switches:
   - 4U in an HA configuration
   - 2U for each NS224 storage shelf
3. Install any required network switches.

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

Step 2: Unpack the boxes

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

Steps

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

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

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

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

Hardware

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

Cables

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

Step 3: Register your storage system

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

Steps

1. Locate the System Serial Numbers (SSN) for every controller being installed. You can find the serial numbers in the following locations:
2. You can find the serial numbers in the following locations:
    - On the packing slip
    - In your confirmation email
    - On each controller's System Management module

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

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

What's next?

After you've prepared to install your AFF A1K hardware, you install the hardware for your AFF A1K storage system.

Install the hardware - AFF A1K

After you prepare to install your AFF A1K storage system, install the hardware for the system. First, install the rail kits. Then install and secure your platform in a cabinet or telco rack.

Skip this step if your cabinet is pre-populated.

Before you begin

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

Steps

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

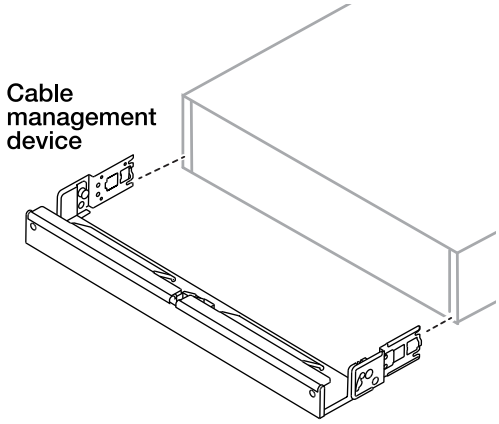

5. Install and secure the shelf as needed.
   a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

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

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

      What's next?

      After you've installed the hardware for your AFF A1K system, you cable the hardware for your AFF A1K storage system.

Cable the hardware - AFF A1K

After you install the rack hardware for your AFF A1K storage system, install the network cables for the controllers, and connect the cables between the controllers and storage shelves.

Before you begin

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

About this task

- These procedures show common configurations. The specific cabling depends on the components ordered for your storage system. For comprehensive configuration and slot priority details, see NetApp Hardware Universe.
- The I/O slots on an AFF A1K controller are numbered 1 through 11.

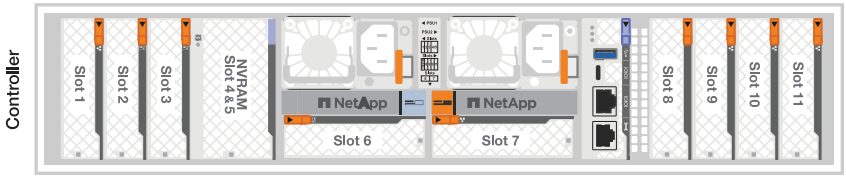

- The cabling graphics have arrow icons showing the proper orientation (up or down) of the cable connector pull-tab when inserting a connector into a port.

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

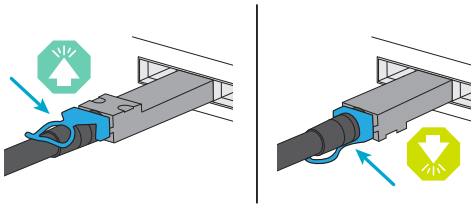

- If cabling to an optical switch, insert the optical transceiver into the controller port before cabling to the switch port.

Step 1: Cable the cluster/HA connections

Cable the controllers to your ONTAP cluster. This procedure differs depending on your storage system model and I/O module configuration.

The cluster interconnect traffic and the HA traffic share the same physical ports.

Switchless cluster cabling

Use the Cluster/HA interconnect cable to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A to port e1a on Controller B.
2. Connect port e7a on Controller A to port e1a on Controller B.

Cluster/HA interconnect cables

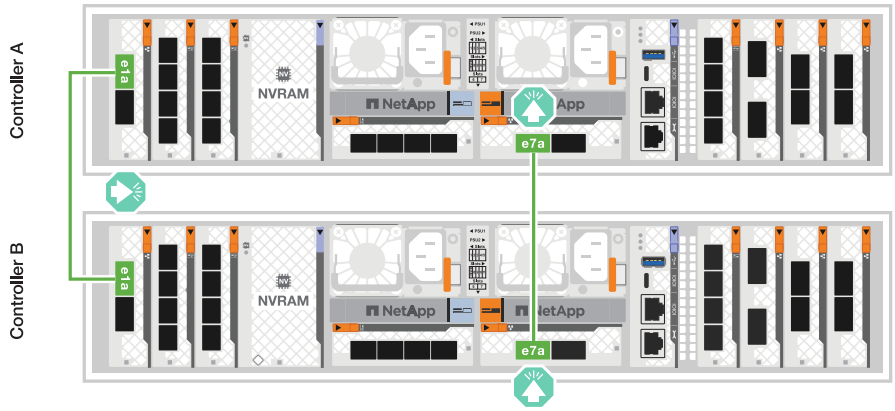

Switched cluster cabling

Use the 100 GbE cable to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A and port e1a on Controller B to cluster network switch A.
2. Connect port e7a on Controller A and port e7a on Controller B to cluster network switch B.

100 GbE cable

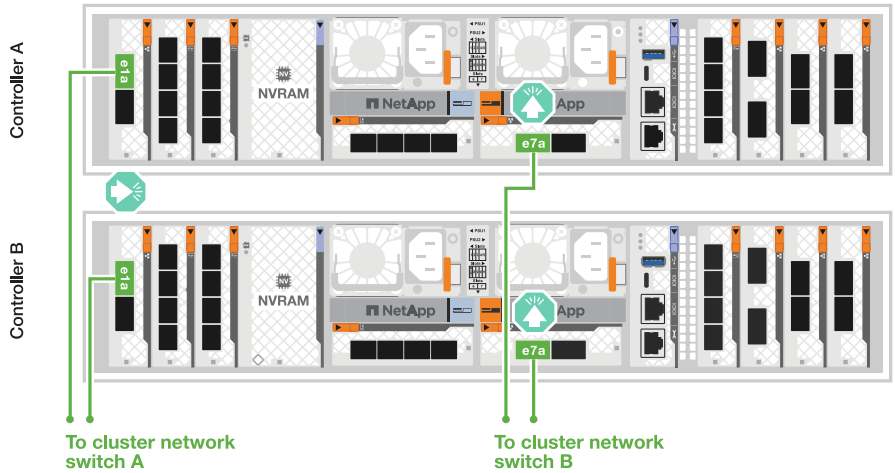

Step 2: Cable the host network connections

Connect the Ethernet module ports to your host network.

The following are some typical host network cabling examples. See NetApp Hardware Universe for your specific system configuration.

Steps

1. Connect ports e9a and e9b to your Ethernet data network switch.

   For maximum system performance for cluster and HA traffic, do not use ports e1b and e7b ports for host network connections. Use a separate host card to maximize performance.

   100 GbE cable

   

   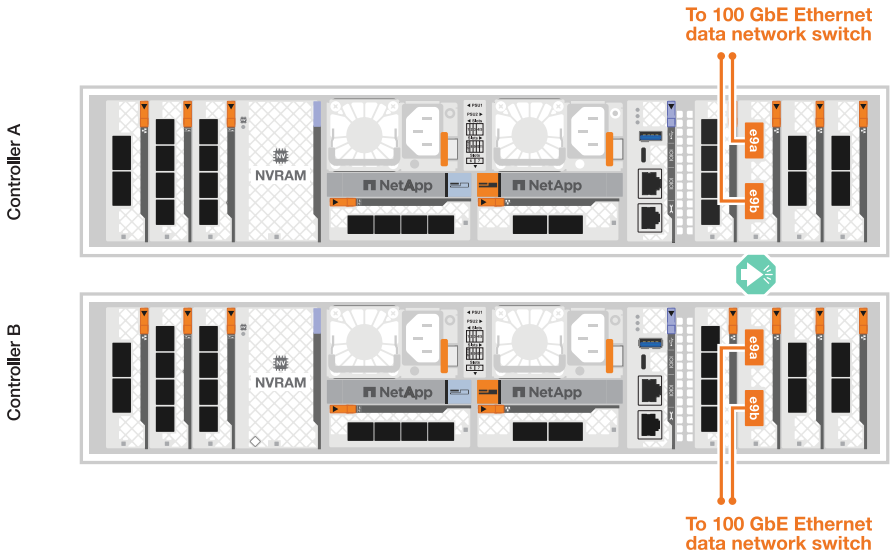

2. Connect your 10/25 GbE host network switches.

   10/25 GbE Host

   

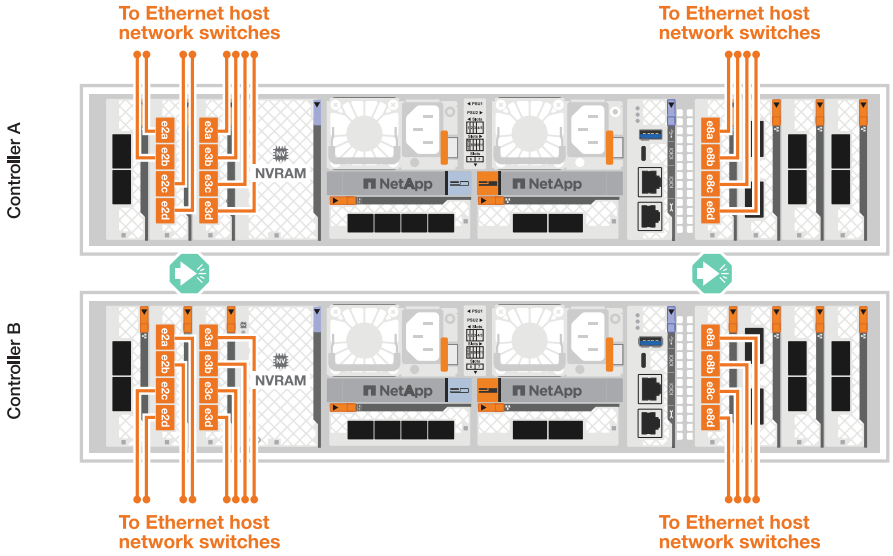

Step 3: Cable the management network connections

Use the 1000BASE-T RJ-45 cables to connect the management (wrench) ports on each controller to the management network switches.

1000BASE-T RJ-45 cables

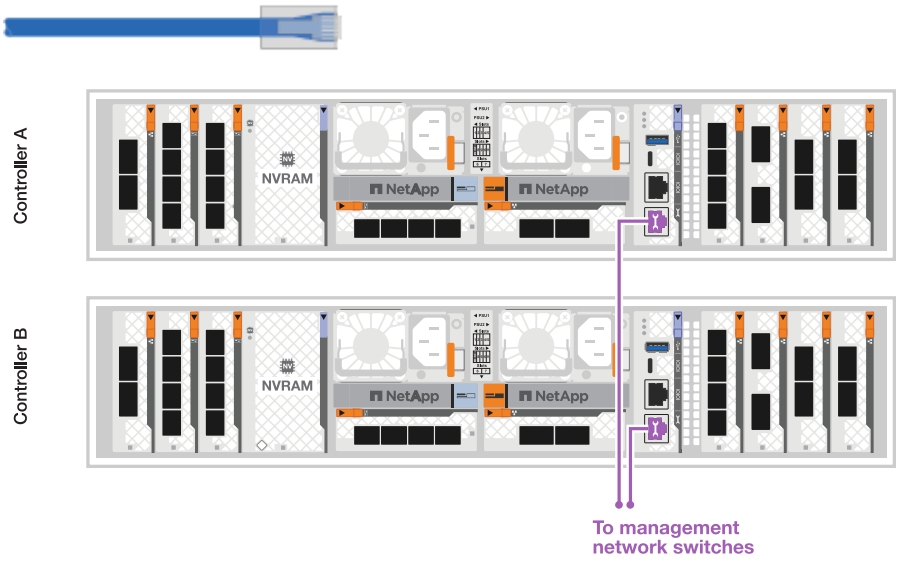

Do not plug in the power cords yet.

Step 4: Cable the shelf connections

The following cabling procedures show how to connect your controllers to a storage shelf. Choose one of the following cabling options that matches your setup.

For the maximum number of shelves supported for your storage system and for all of your cabling options, see NetApp Hardware Universe.

About this task

The AFF A1K storage systems support NS224 shelves with either the NSM100 or NSM100B module. The

major differences between the modules are:

- NSM100 shelf modules use built-in port e0a and e0b.
- NSM100B shelf modules use ports e1a and e1b in slot 1.

The following cabling example shows NSM100 modules in the NS224 shelves when referring to shelf module ports.

Option 1: Connect to one NS224 storage shelf

Connect each controller to the NSM modules on the NS224 shelf. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables

Steps

1. On controller A, connect the following ports:
   a. Connect port e11a to NSM A port e0a.
   b. Connect port e11b to port NSM B port e0b.

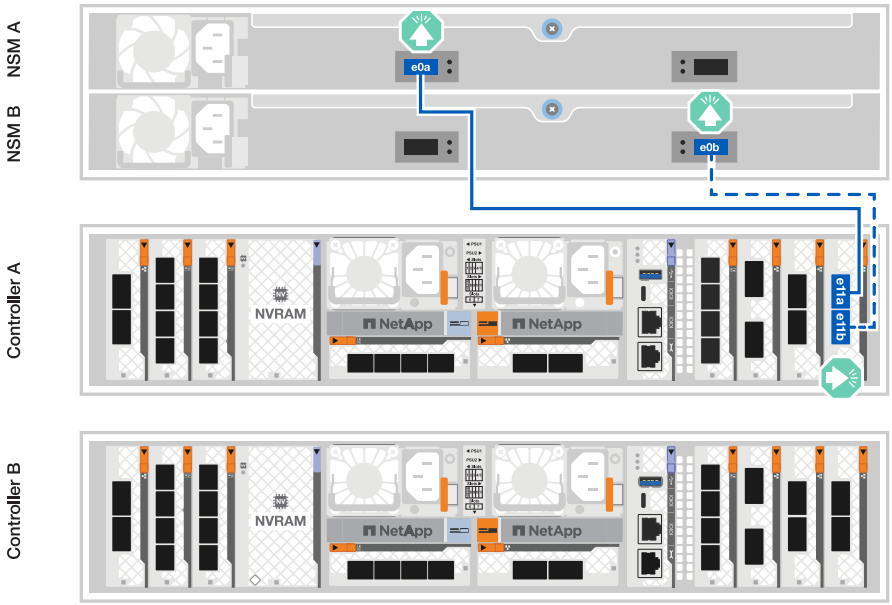

2. On controller B, connect the following ports:
   a. Connect port e11a to NSM B port e0a.
   b. Connect port e11b to NSM A port e0b.

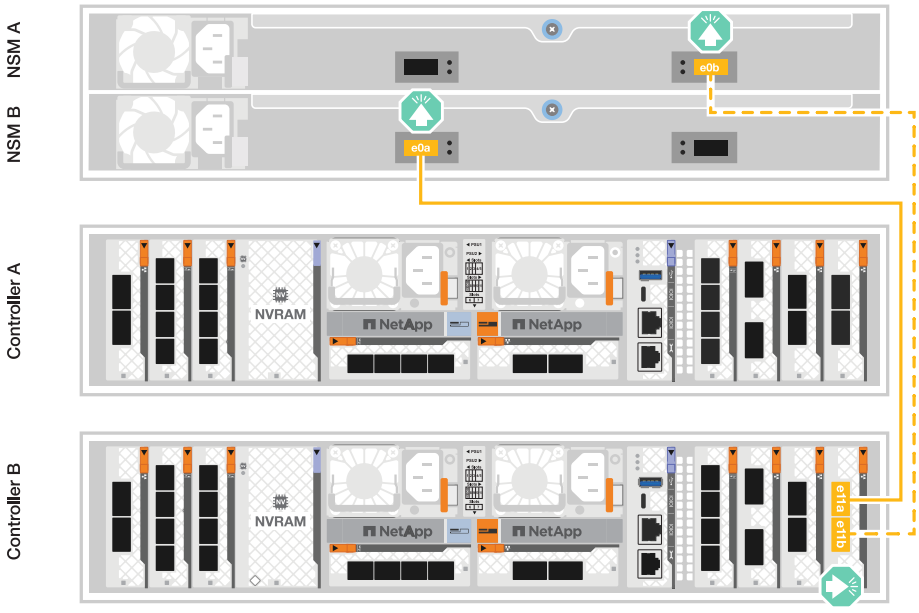

Option 2: Connect to two NS224 storage shelves

Connect each controller to the NSM modules on both NS224 shelves. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables

Steps

1. On controller A, connect the following ports:
   a. Connect port e11a to shelf 1 NSM A port e0a.
   b. Connect port e11b to shelf 2 NSM B port e0b.
   c. Connect port e10a to shelf 2 NSM A port e0a.
   d. Connect port e10b to shelf 1 NSM A port e0b.

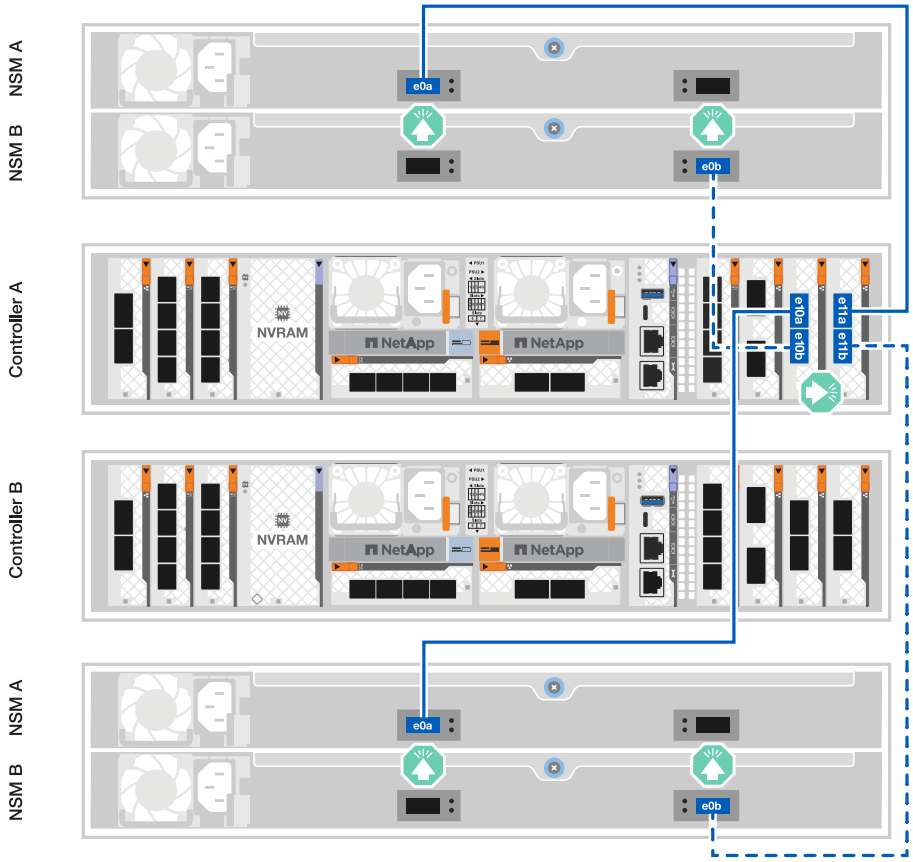

2. On controller B, connect the following ports:
   a. Connect port e11a to shelf 1 NSM B port e0a.
   b. Connect port e11b to shelf 2 NSM A port e0b.
   c. Connect port e10a to shelf 2 NSM B port e0a.
   d. Connect port e10b to shelf 1 NSM A port e0b.

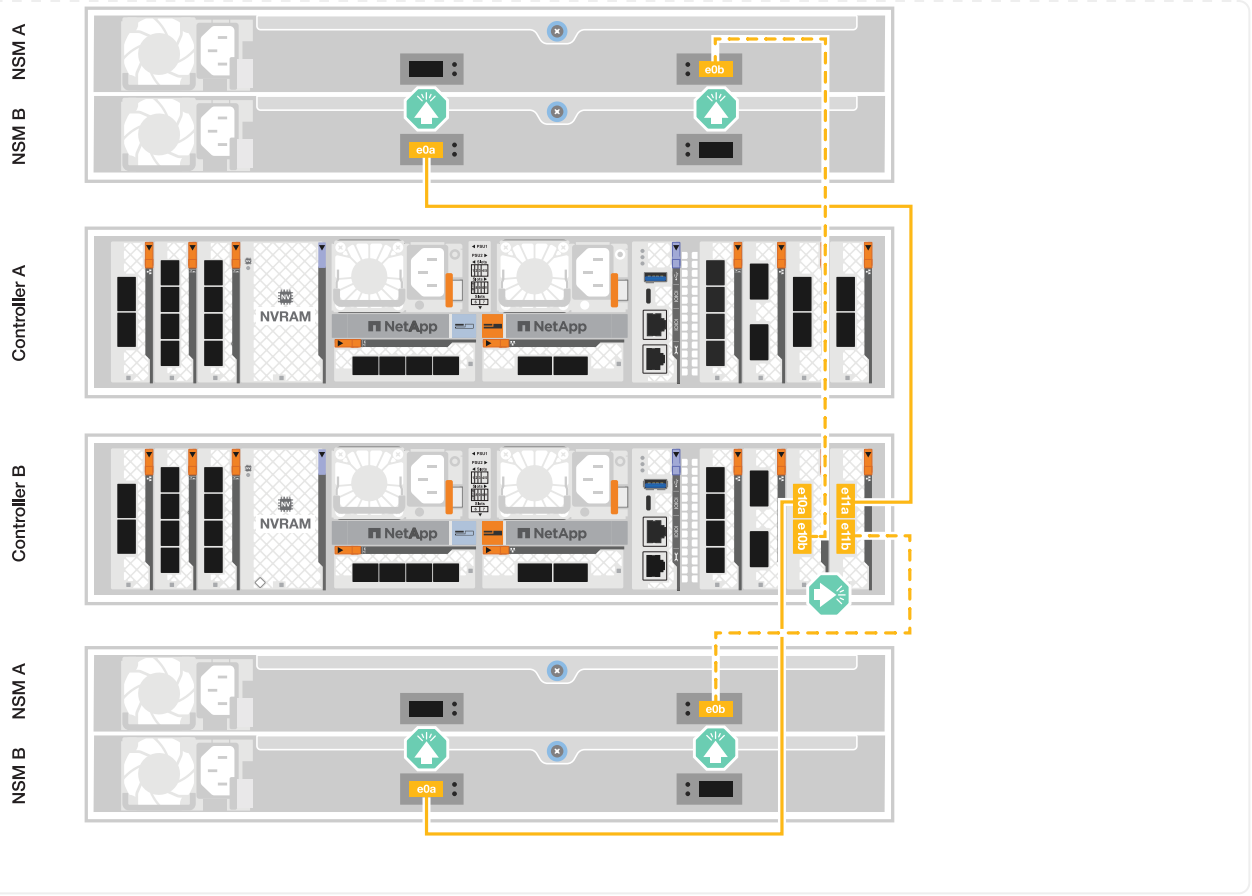

What's next?

After you've cabled the hardware for your AFF A1K system, you power on the AFF A1K storage system.

Power on the storage system - AFF A1K

After you install the rack hardware for your AFF A1K storage system and install the cables for the controllers and storage shelves, you should power on your storage shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

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

Before you begin

Make sure you have a paperclip or narrow tipped ball point pen for setting NS224 storage shelf IDs.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed shelf ID of 00.

- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, and then plug them back in) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, and then connecting the power cords to power sources on different circuits.

   The shelf powers on and boots automatically when plugged into the power source.

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

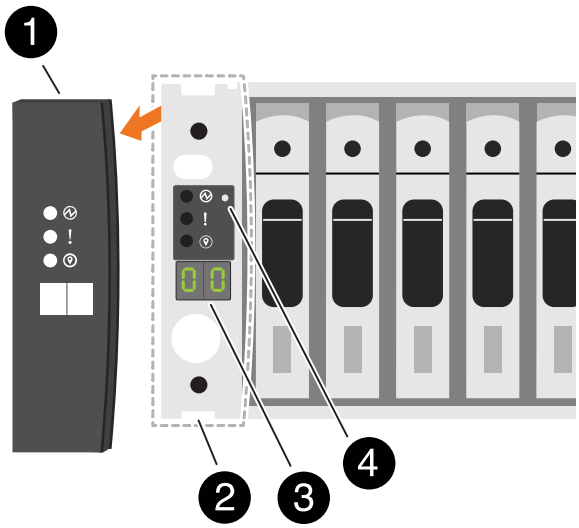

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID number
4	Shelf ID button

3. Change the first number of the shelf ID:
   a. Insert the straightened end of a paperclip or narrow tipped ball point pen into the small hole to press the shelf ID button.
   b. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

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

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

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

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

      The first number continues to blink.

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

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

      The first number on the digital display stops blinking.

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

      The second number continues to blink.

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

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

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

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.
   a. Unplug the power cord from both power supplies on the shelf.
   b. Wait 10 seconds.
   c. Plug the power cords back into the shelf power supplies to complete the power cycle.

      A power supply is powered on as soon as the power cord is plugged in. Its bicolored LED should illuminate green.

7. Replace the left end cap.

Step 2: Power on the controllers

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

Steps

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

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

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

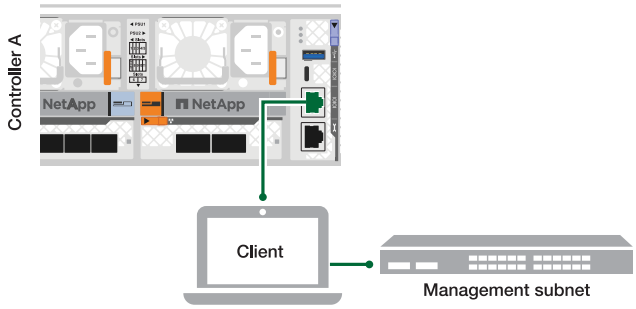

2. Assign a TCP/IP address to the laptop, using one that is on the management subnet.
3. Plug the two power cords into the controller power supplies, and then connect them to power sources on different circuits.

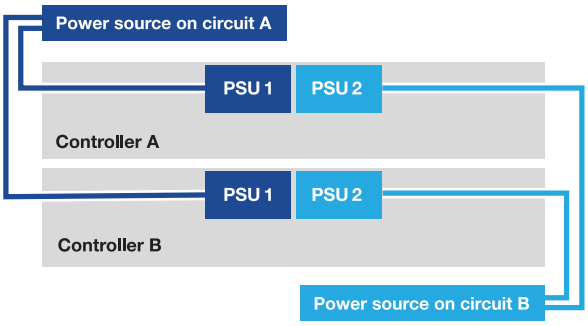

   - The system begins to boot. Initial booting might take up to eight minutes.
   - The LEDs flash on and the fans start, which indicates that the controllers are powering on.
   - The fans might be very noisy when they first start up. The fan noise during start-up is normal.
4. Secure the power cords using the securing device on each power supply.

What's next?

After you've turned on your AFF A1K storage system, you set up your cluster.

Maintain

Overview of the maintenance procedures - AFF A1K

Maintain the hardware of your AFF A1K storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the AFF A1K system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF A1K storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media- manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the system uses to boot the image from a USB drive and restore the configuration from the partner node.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Fan	A fan cools the controller.
NVRAM	The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module.
NV battery	The NV battery is responsible for providing power to the NVRAM module while data in-flight is being destaged to flash memory after a power loss.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.
System management module	The System management module provides the interface between the controller and a console or laptop for controller or system maintenance purposes. The System management module contains the boot media and stores the system serial number (SSN).

Boot media - automated recovery

Boot media automated recovery workflow - AFF A1K

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node

to reinstall ONTAP on the replacement boot media in your AFF A1K storage system.

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

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

Review the boot media requirements

Review the requirements for boot media replacement.

Shut down the controller

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

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media.

Restore the image on the boot media

Restore the ONTAP image from the partner controller.

Return the failed part to NetApp

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

Requirements for automated boot media recovery - AFF A1K

Before replacing the boot media in your AFF A1K system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming the cluster ports on the impaired controller are working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

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

Review the following requirements.

- You must replace the failed boot media with a replacement boot media you received from NetApp.
- The cluster ports are used to communicate between the two controllers during the automated boot recovery process. Make sure that the cluster ports on the impaired controller are working properly.
- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg

- /cfcard/kmip/certs/client.crt
- /cfcard/kmip/certs/client.key
- /cfcard/kmip/certs/CA.pem

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

What's next

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

Shut down the controller for automated boot media recovery - AFF A1K

Shut down the impaired controller in your AFF A1K storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

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

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

About this task

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

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

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

Steps

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

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

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

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

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

What's next

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

Replace the boot media for automated boot recovery - AFF A1K

The boot media in your AFF A1K system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media in the System Management module, and then reinstalling the System Management module.

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

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

Replace the boot media.

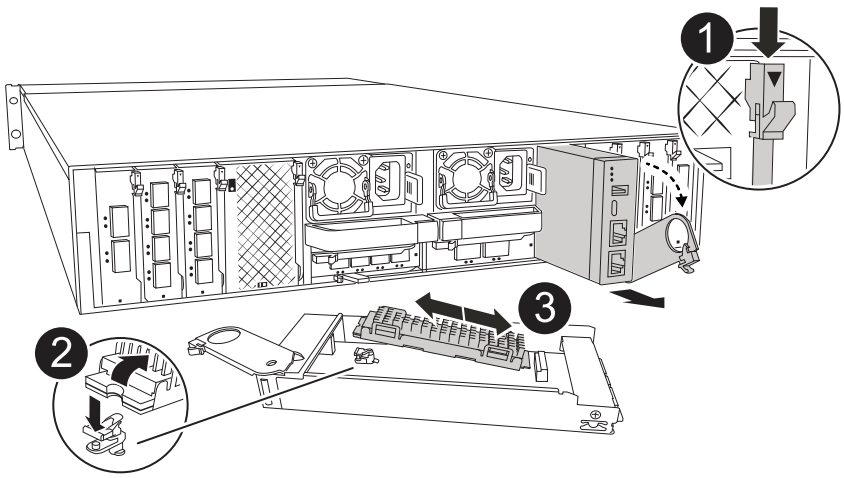

1	System Management module cam latch

2	Boot media locking button
3	Boot media

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.

   If your storage system has DC power supplies, disconnect the power cable block from the power supply units (PSUs).

3. Remove the System Management module:
   a. Remove any cables connected to the System Management module. Make sure that you label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   c. Depress the System Management cam button.
   d. Rotate the cam latch down as far as it will go.
   e. Remove the System Management module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
   f. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:
   a. Press the blue locking button.
   b. Rotate the boot media up, slide it out of the socket, and set it aside.
5. Install the replacement boot media into the System Management module:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the locking button.
   c. Push the locking button, rotate the boot media all the way down and then release the locking button.
6. Reinstall the System Management module:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
7. Rotate the cable management tray up to the closed position.
   a. Recable the System Management module.
8. Plug the power cables into the power supplies and reinstall the power cable retainer.

   The controller begins to boot as soon as power is reconnected to the system.

What's next

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

Automated boot media recovery from the partner node - AFF A1K

After installing the new boot media device in your AFF A1K system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

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

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

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

2. Monitor the boot media install recovery process.

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

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

   Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

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

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

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

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

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

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

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

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

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

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

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

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

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

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

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

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

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

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

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

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

5. If automatic giveback was disabled, reenable it:

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

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

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

What's next

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

Return the failed boot media part to NetApp - AFF A1K

If a component in your AFF A1K system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF A1K

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the AFF A1K system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

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

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

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

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

Shut down the controller

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

Replace the boot media

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

Restore the image on the boot media

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

6 Restore encryption

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

7 Return the failed part to NetApp

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

Requirements for manual boot media recovery - AFF A1K

Before replacing the boot media in your AFF A1K system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

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

File preparation

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

Component replacement

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

Controller identification

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

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

What's next?

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

Check encryption support for manual boot media recovery - AFF A1K

To ensure data security on your AFF A1k storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

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

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

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

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

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

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

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

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

What's next?

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

Shut down the controller for manual boot media recovery - AFF A1K

Shut down the impaired controller in your AFF A1K storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

What's next?

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

Replace the boot media and prepare for manual boot recovery - AFF A1K

The boot media in your AFF A1K system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

Step 1: Replace the boot media

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

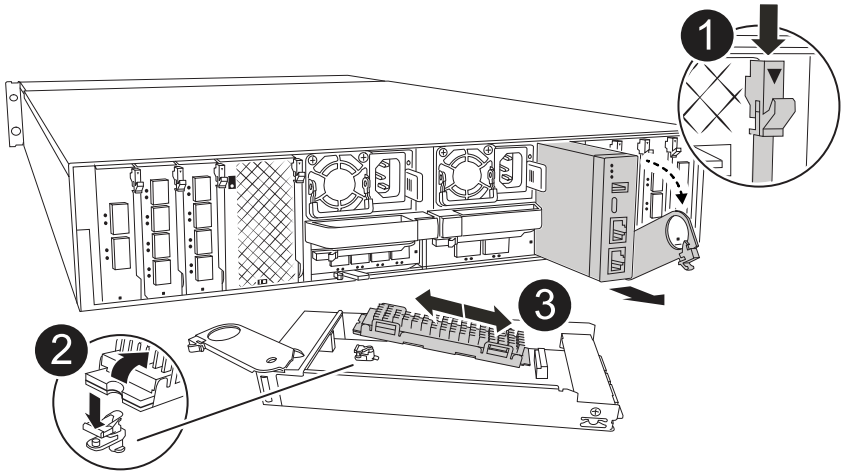

1	System Management module cam latch
2	Boot media locking button
3	Boot media

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.

   If your storage system has DC power supplies, disconnect the power cable block from the power supply units (PSUs).

3. Remove the System Management module:
   a. Remove any cables connected to the System Management module. Make sure that you label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   c. Depress the System Management cam button.
   d. Rotate the cam latch down as far as it will go.
   e. Remove the System Management module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
   f. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:
   a. Press the blue locking button.
   b. Rotate the boot media up, slide it out of the socket, and set it aside.
5. Install the replacement boot media into the System Management module:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the locking button.
   c. Push the locking button, rotate the boot media all the way down and then release the locking button.
6. Reinstall the System Management module.
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
7. Rotate the cable management tray up to the closed position.
   a. Recable the System Management module.

Step 2: Transfer the ONTAP image to the boot media

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

Before you begin

- You must have an empty USB flash drive, formatted to FAT32, with at least 4GB capacity.
- Download a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site. Use the `version -v` command to display if your version of ONTAP supports NVE. If the command output displays

`<10no- DARE>`, your version of ONTAP does not support NVE.

  - If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

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

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

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

   c. Remove the USB flash drive from your laptop.
2. Insert the USB flash drive into the USB slot on the System Management module.

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

3. Plug the power cables back into the PSU.

   The system will begin to reboot, typically to the LOADER prompt.

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

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

What's next?

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

Manual boot media recovery from a USB drive - AFF A1K

After installing the new boot media device in your AFF A1K system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Steps

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

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.

3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

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

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

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

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.

5. Give back the controller using the `storage failover giveback -fromnode local` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

   NOTE: If the process fails, contact NetApp Support.

What's next?

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

Restore encryption keys after manual boot recovery - AFF A1K

Restore encryption on the replacement boot media in your AFF A1K system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

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

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

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

3. Confirm that you want to continue the recovery process.

   Show example prompt

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

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

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

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

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

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

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

Successfully recovered keymanager secrets.

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

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

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

Show example prompt

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


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

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

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

8. From the partner node, giveback the partner controller by entering the following command.

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

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

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

Show example prompt

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

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

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

11. Ensure that all keys are synced by entering the following command.

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

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

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

13. Restore automatic giveback, if you disabled it, by entering the following command.

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

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

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

Option 2: Restore the External Key Manager configuration

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

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

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

   Show example boot menu

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

3. When prompted, confirm you have gathered the required information.

   Show example prompt

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

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

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

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

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

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

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

After you enter the client and server information, the recovery process completes.

Show example

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

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

Show example prompt

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

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

6. Restore automatic giveback if you disabled it.

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

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

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

What's next?

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

Return the failed boot media part to NetApp - AFF A1K

If a component in your AFF A1K system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - AFF A1K

Get started with replacing the controller in your AFF A1K storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

1 Review the requirements to replace the controller

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

2 Shut down the impaired controller

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

3 Replace the controller

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

4 Restore and verify the system configuration

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

5 Give back the controller

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

6 Complete controller replacement

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

Requirements to replace the controller - AFF A1K

Before replacing the controller in your AFF A1K system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements for replacing the controller.

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

- You cannot change any drives or drive shelves as part of this procedure.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text log file.

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

What's next?

After you've the reviewed the requirements to replace your AFF A1K controller, you need to shut down the impaired controller.

Shut down the impaired controller - AFF A1K

Shut down the controller in your AFF A1K storage system to prevent data loss and ensure system stability when replacing the controller.

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

What's next?

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

Replace the controller - AFF A1K

Replace the controller in your AFF A1K system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

Step 1: Remove the controller module

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

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

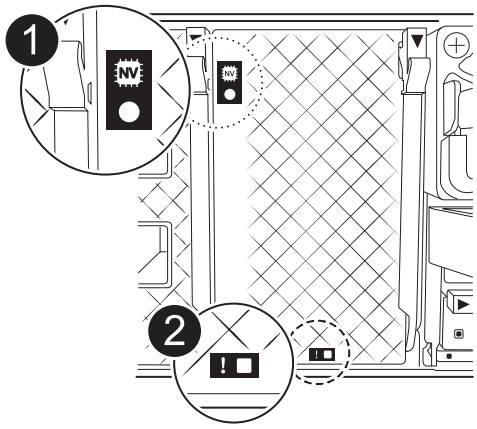

1	NVRAM status LED
2	NVRAM attention LED

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

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

2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

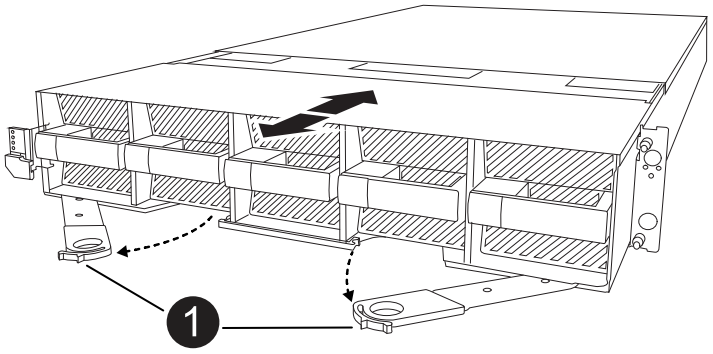

1	Locking cam latches

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

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

Step 2: Move the fans

You must remove the five fan modules from the impaired controller module to the replacement controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press the gray locking button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

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

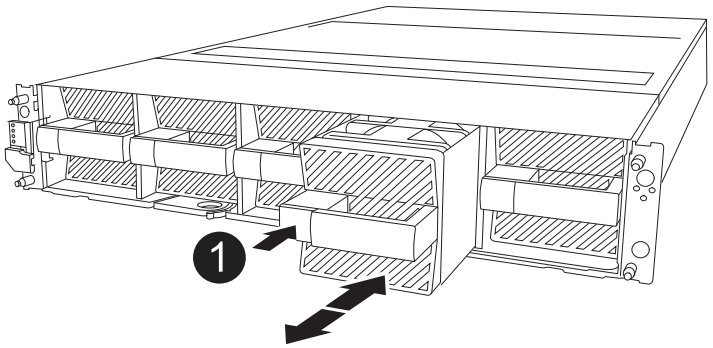

1	Black locking button

4. Install the fan in the replacement controller module:
   a. Align the edges of the fan housing with the opening in the front of the replacement controller module.
   b. Gently slide the fan module all the way into the replacement controller module until it locks in place.
5. Repeat the preceding steps for the remaining fan modules.

Step 3: Move the NV battery

Move the NV battery to the replacement controller.

Steps

1. Open the NV battery air duct cover and locate the NV battery.

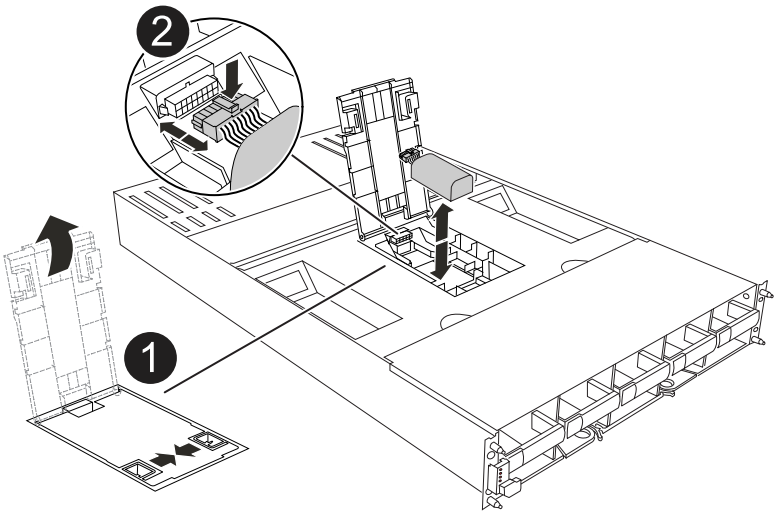

1	NV battery air duct cover
2	NV battery plug
3	NV battery pack

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module.
5. Move the battery pack to the replacement controller module and then install it in the NV battery air duct:
   a. Open the NV battery air duct in the replacement controller module.
   b. Plug the battery plug into the socket and make sure that the plug locks into place.
   c. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.

d. Close the air duct cover.

Step 4: Move system DIMMs

Move the DIMMs to the replacement controller module.

Steps

1. Open the motherboard air duct and locate the DIMMs.

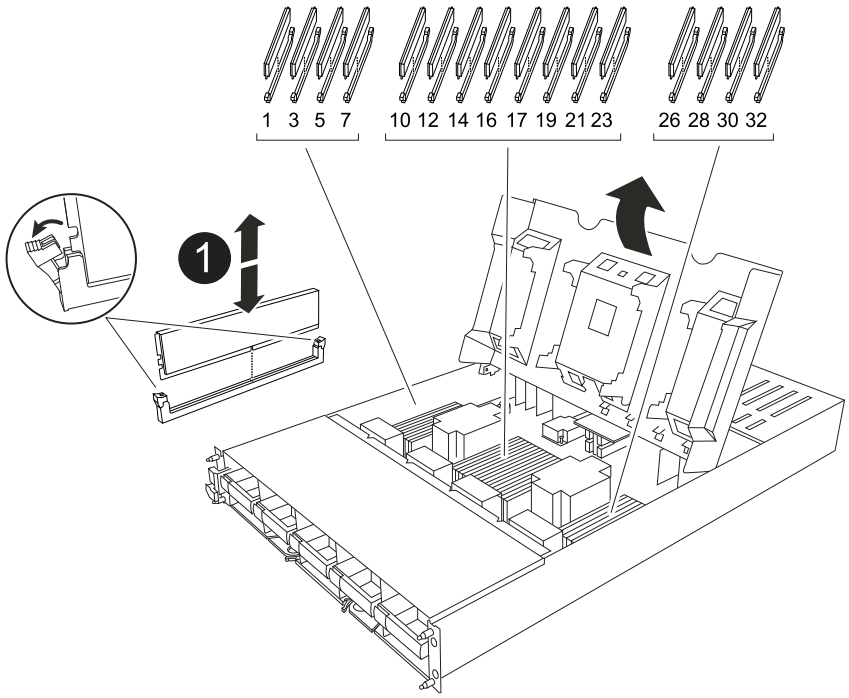

1	System DIMM

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

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

4. Locate the slot where you are installing the DIMM in the replacement controller module.
5. Insert the DIMM squarely into the slot.

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

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

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

7. Repeat these steps for the remaining DIMMs.
   Close the motherboard air duct.

Step 5: Install the controller module

Reinstall the controller module and boot it.

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

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

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

   The controller boots to the LOADER prompt as soon as it is fully seated.

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

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

5. Set the current time in GMT with the `set time hh:mm:ss` command. You can get the current GMT from the partner node the `date -u`command.
6. Recable the storage system, as needed.

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

What's next?

After you've replaced the impaired AFF A1K controller, you need to restore the system configuration.

Restore and verify the system configuration - AFF A1K

Verify that the controller's HA configuration is active and functioning correctly in your AFF A1K storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

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

Steps

1. Boot to maintenance mode: `boot_ontap maint`
   a. Enter `y` when you see *Continue with boot?*.

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

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

   If you see *PERSONALITY MISMATCH* contact customer support.

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

   The HA state should be the same for all components.

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

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

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

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

Step 2: Verify disk list

Steps

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

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

2. Exit Maintenance mode: `halt`.

What's next?

After you've restored and verified the system configuration for your AFF A1K system, you need to give back the controller.

Give back the controller - AFF A1K

Return control of storage resources to the replacement controller so your AFF A1K system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

No encryption

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

Steps

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

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

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

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`
   - If you encounter errors, contact NetApp Support.
9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

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

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

The output should show no results when when properly synchronized.

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

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

What's next?

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

Complete controller replacement - AFF A1K

To complete the controller replacement for your AFF A1K system, first restore the NetApp Storage Encryption configuration (if necessary). Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, return the failed part to NetApp.

Step 1: Verify LIFs and check cluster health

Before returning the replacement node to service, verify that the logical interfaces are on their home ports, check the cluster health, and reset automatic giveback.

Steps

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

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

2. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 2: Return the failed part to NetApp

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

Replace a DIMM - AFF A1K

Replace a DIMM in your AFF A1K system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting

ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Remove the controller module

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

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

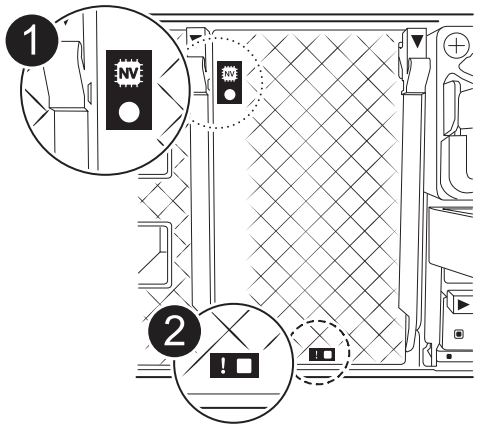

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

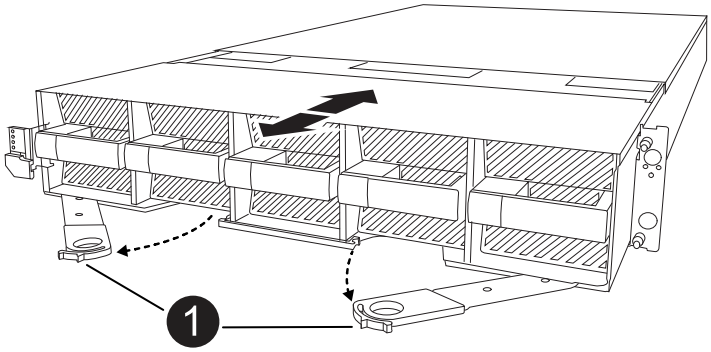

1	Locking cam latches

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

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

Step 3: Replace a DIMM

You must replace a DIMM when the system reports a permanent failure condition for that DIMM.

Steps

1. If you are not already grounded, properly ground yourself.
2. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
3. Locate the DIMMs on your controller module and identify the DIMM for replacement.

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

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

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

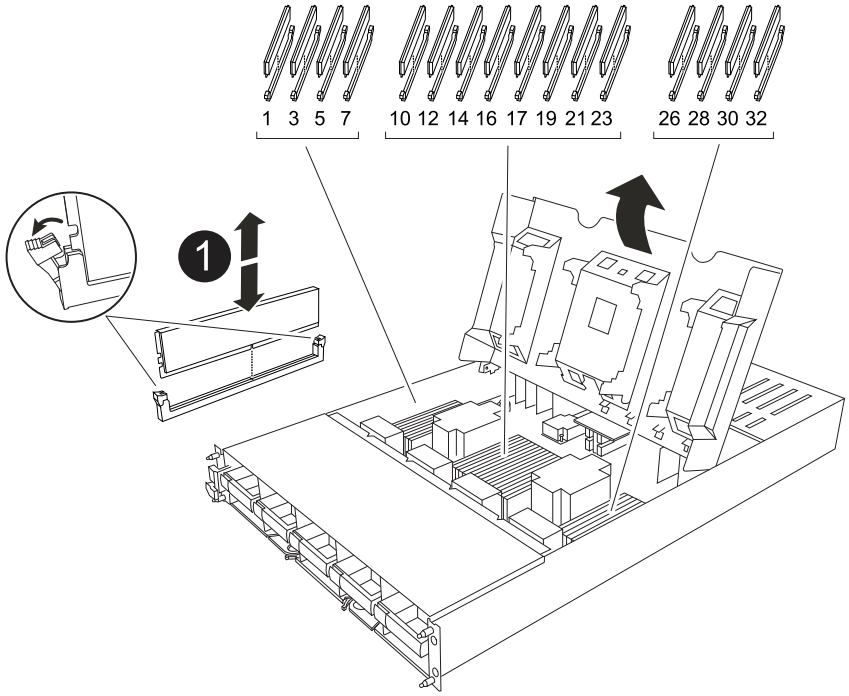

1	DIMM and DIMM ejector tabs

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

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

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

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

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

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

Step 4: Install the controller

Reinstall the controller module and boot it.

Steps

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

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.

3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

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

   The controller module begins to boot as soon as it is fully seated in the enclosure.

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

Step 5: Return the failed part to NetApp

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

Replace a fan - AFF A1K

Replace a fan module in your AFF A1K system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

Steps

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

   Facing the controller module, fan modules are numbered 1 through 5, from left to right.

   There is a single LED for each fan. It is green when the fan is functioning correctly and amber when not.

4. Press the black button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

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

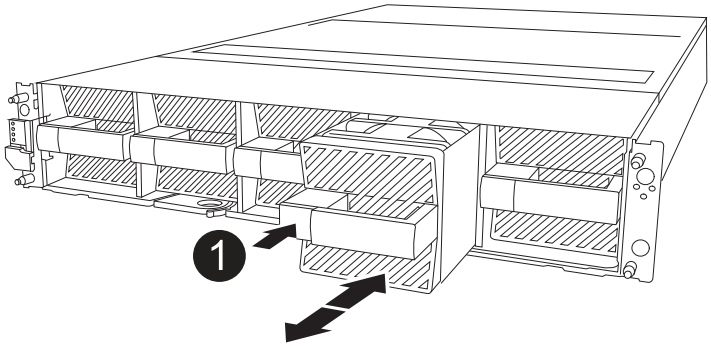

1	Black release button

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED turns off once the fan is recognized by that system.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace NVRAM - AFF A1K

Replace the NVRAM in your AFF A1K system when the non-volatile memory becomes faulty or requires an upgrade. The replacement process involves shutting down the impaired controller, replacing the NVRAM module or the NVRAM DIMM, reassigning the disks, and returning the failed part to NetApp.

The NVRAM module consists of the NVRAM12 hardware and field-replaceable DIMMs. You can replace a failed NVRAM module or the DIMMs inside the NVRAM module.

Before you begin

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Replace the NVRAM module or NVRAM DIMM

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

Option 1: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 4/5 in the enclosure and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
4. Remove the impaired NVRAM module from the enclosure:
   a. Depress the locking cam button.

      The cam button moves away from the enclosure.

   b. Rotate the cam latch down as far as it will go.
   c. Remove the impaired NVRAM module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

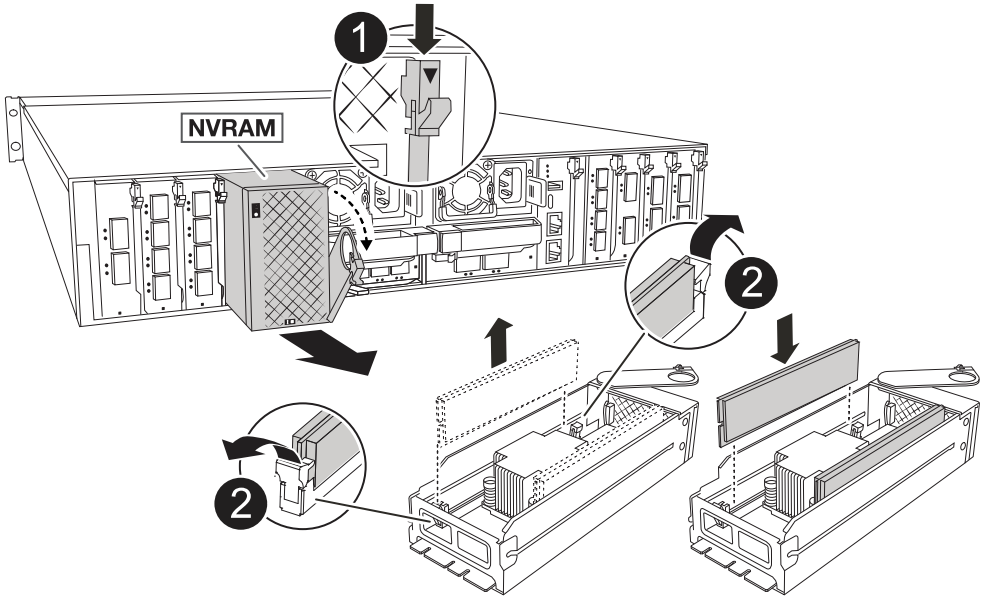

1	Cam locking button
2	DIMM locking tabs

5. Set the NVRAM module on a stable surface.
6. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
7. Install the replacement NVRAM module into the enclosure:
   a. Align the module with the edges of the enclosure opening in slot 4/5.

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

8. Recable the controller.
9. Rotate the cable management tray up to the closed position.

Option 2: Replace the NVRAM DIMM

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

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
4. Remove the target NVRAM module from the enclosure.

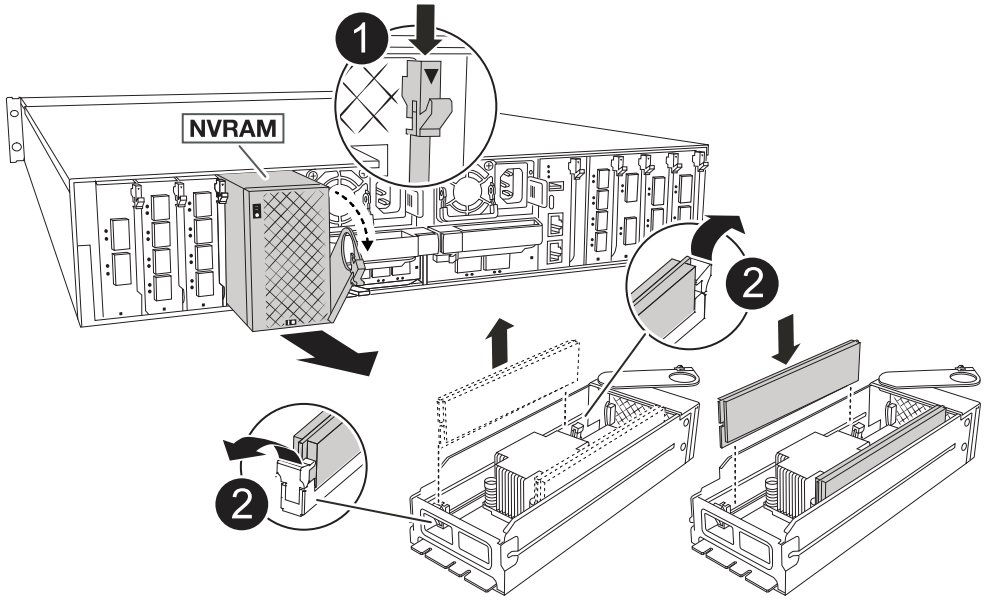

1	Cam locking button
2	DIMM locking tabs

5. Set the NVRAM module on a stable surface.
6. Locate the DIMM to be replaced inside the NVRAM module.

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

7. Remove the DIMM by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
8. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.

9. Install the NVRAM module into the enclosure:
    a. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.
10. Recable the controller.
11. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace the FRU, you must reboot the controller module by plugging the power cables back into the PSU.

Steps

1. Plug the power cables back into the PSU.

   The system will begin to reboot, typically to the LOADER prompt.

2. Enter *bye* at the LOADER prompt.
3. Return the impaired controller to normal operation by giving back its storage: *`storage failover giveback -ofnode _impaired_node_name`*.
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true .`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 4: Reassign disks

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

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

Steps

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

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

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

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

      The controller takes back its storage and completes booting.

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

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

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

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

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

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

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

```
node1:> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        151759706  151759706  -
151759706 Pool0
1.0.1  aggr0_1  node1 node1           151759706  151759706  -
151759706 Pool0
.
.
.
```

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

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

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

   This is required if both of the following are true:

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

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

8. If your system is in a MetroCluster configuration, verify that each controller is configured: *metrocluster node show - fields configuration-state*

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

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

4 entries were displayed.
```

9. Verify that the expected volumes are present for each controller: `vol show -node node-name`
10. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.
11. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
12. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

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

Replace the NV battery - AFF A1K

Replace the NV battery in your AFF A1K system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Remove the controller module

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

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

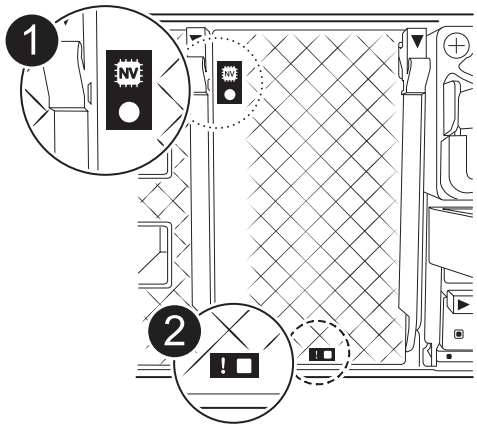

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

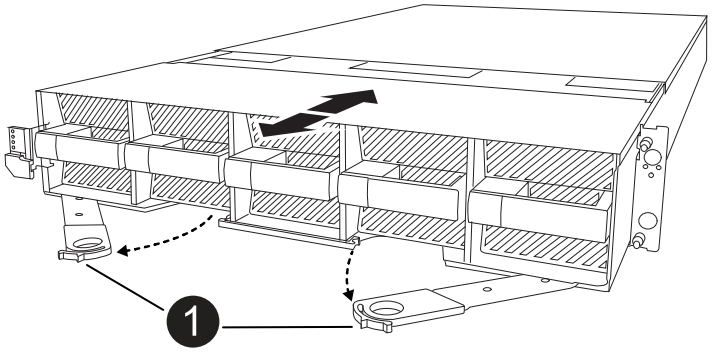

1	Locking cam latches

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

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

Step 3: Replace the NV battery

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

Steps

1. Open the air duct cover and locate the NV battery.

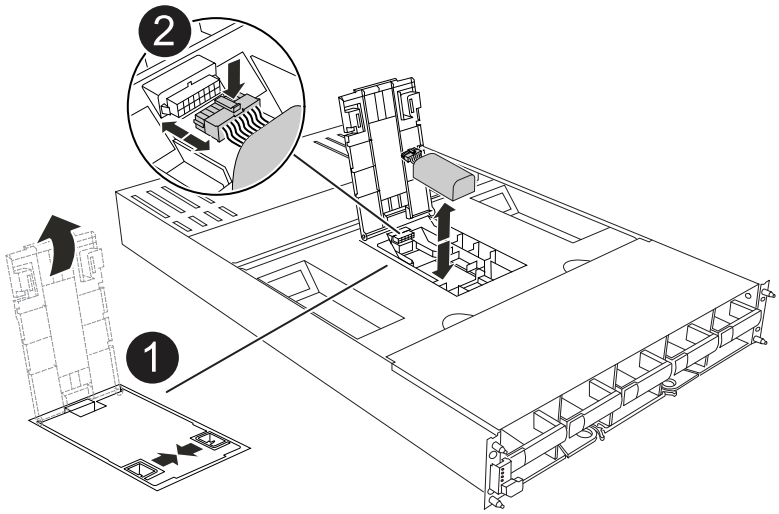

1	NV battery air duct cover
2	NV battery plug

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module, and then set it aside.

5. Remove the replacement battery from its package.
6. Install the replacement battery pack into the controller:
   a. Plug the battery plug into the riser socket and make sure that the plug locks into place.
   b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
7. Close the NV air duct cover.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

Steps

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

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

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

   The controller module begins to boot as soon as it is fully seated in the enclosure.

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

Step 5: Return the failed part to NetApp

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

I/O module

Overview of add and replace an I/O module - AFF A1K

The AFF A1K system offers flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your AFF A1K storage system with the same type of I/O module, or with

a different kind of I/O module. You can also add an I/O module into a system with empty slots.

- Add an I/O module

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

- Replace an I/O module

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

I/O slot numbering

The I/O slots on the AFF A1K controller are numbered 1 through 11, as shown in the following illustration.

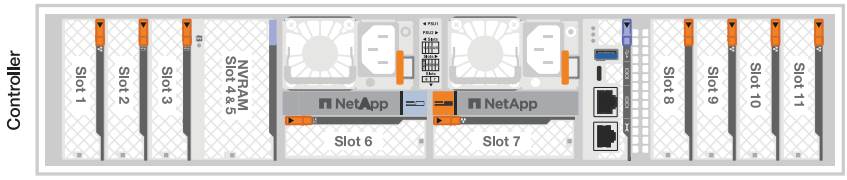

Add an I/O module - AFF A1K

Add an I/O module to your AFF A1K system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your AFF A1K storage system when there are empty slots available or when all slots are fully populated.

About this task

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

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

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

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: MetroCluster configuration

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

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

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

Steps

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

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

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

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

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Add the new I/O module

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

Before you begin

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

Add I/O module to an available slot

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

Steps

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the target slot blanking module from the carrier:
   a. Depress the cam latch on the blanking module in the target slot.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
4. Install the I/O module:
   a. Align the I/O module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
5. Cable the I/O module to the designated device.

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

6. Rotate the cable management tray up to the closed position.
7. From the LOADER prompt, reboot the node:

   ```
   bye
   ```

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

8. Give back the controller from the partner controller:

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

9. Repeat these steps for controller B.
10. From the healthy node, restore automatic giveback if you disabled it:

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

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

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

Add I/O module to a fully-populated system

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

About this task

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

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

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the target I/O module from the chassis:
   a. Depress the cam latch button.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

      Make sure that you keep track of which slot the I/O module was in.
5. Install the I/O module into the target slot in the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
6. Cable the I/O module to the designated device.
7. Repeat the remove and install steps to replace additional modules for the controller.
8. Rotate the cable management tray up to the closed position.
9. Reboot the controller from the LOADER prompt:_bye_

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

   If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement
10. Give back the controller from the partner controller:

    ```
    storage failover giveback -ofnode target_node_name
    ```
11. Enable automatic giveback if it was disabled:

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

12. Do one of the following:
    - If you removed a NIC I/O module and installed a new NIC I/O module, use the following network command for each port:

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

    - If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
13. Repeat these steps for controller B.

Replace an I/O module - AFF A1K

Replace an I/O module in your AFF A1K system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

You can use this procedure with all versions of ONTAP supported by your storage system.

Before you begin

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

Step 1: Shut down the impaired node

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

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

Steps

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

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

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

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

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace a failed I/O module

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

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

   This following illustration shows removing a horizontal and vertical I/O module. Typically, you will only remove one I/O module.

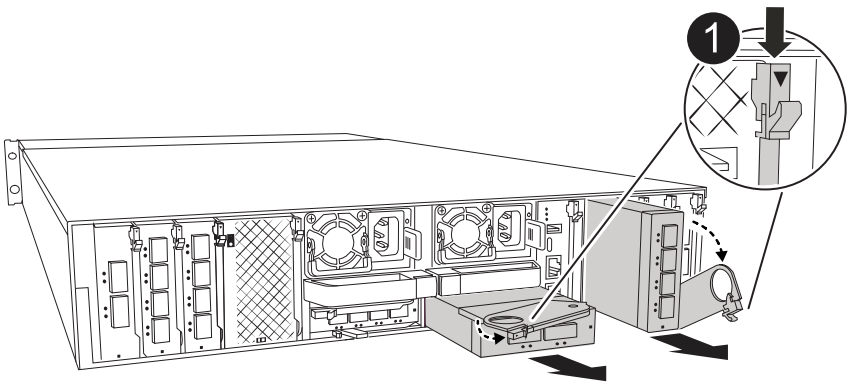

1	I/O cam latch

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

4. Remove the target I/O module from the enclosure:
   a. Depress the cam button on the target module.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

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

5. Set the I/O module aside.
6. Install the replacement I/O module into the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
7. Cable the I/O module.
8. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

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

Steps

1. Reboot the controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

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

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

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

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

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

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

Step 4: Return the failed part to NetApp

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

Replace a power supply - AFF A1K

Replace an AC or DC power supply unit (PSU) in your AFF A1K system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cable, replacing the faulty PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

About this task

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

Do not mix PSUs with different efficiency ratings. Always replace like for like.

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

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

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

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

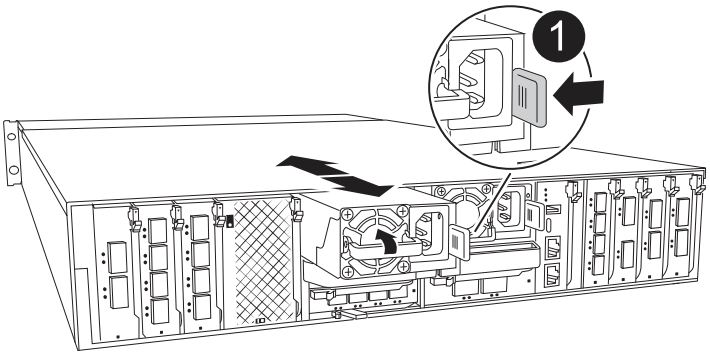

1	Terracotta PSU locking tab

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

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

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

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

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

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

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

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

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

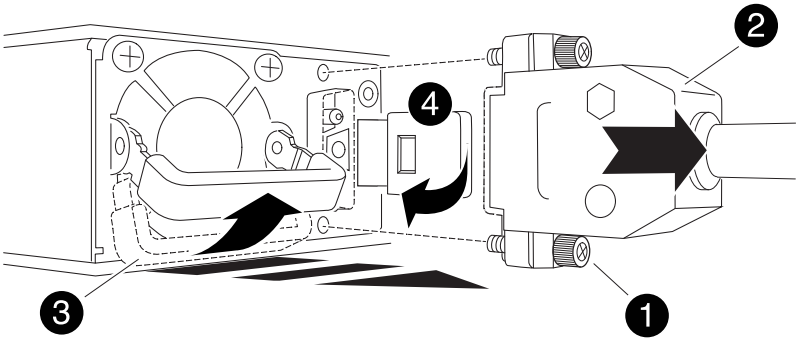

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle
4	Blue PSU locking tab

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

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

way.

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

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A1K

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your AFF A1K system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Remove the controller module

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

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

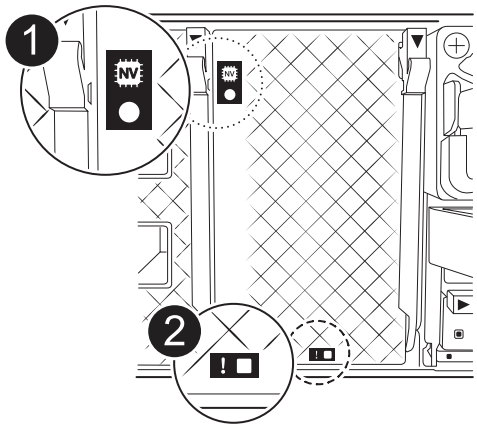

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

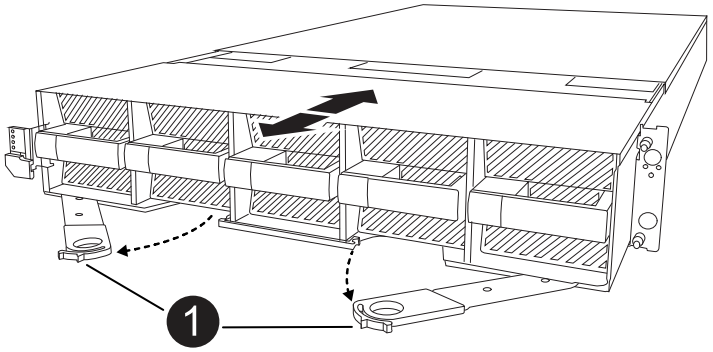

1	Locking cam latches

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

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

Step 3: Replace the RTC battery

Remove failed RTC battery and install the replacement RTC battery.

Steps

1. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the RTC battery under the air duct.

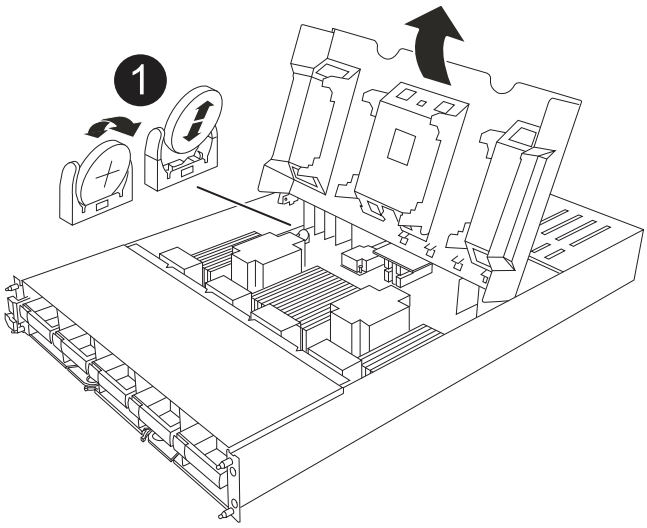

1	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

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

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

Steps

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

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

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

   The controller module begins to boot as soon as it is fully seated in the enclosure.

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

Step 5: Reset the time and date on the controller

After replacing the RTC battery, inserting controller and powering on first BIOS reset, you will see the following error messages:
`RTC date/time error. Reset date/time to default`
`RTC power failure error`
These messages are excpected and you can continue with this procedure.

Steps

1. Check the date and time on the healthy controller with the *cluster date show* command.

If your system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*

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

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

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

   1. Confirm the date and time on the target controller.
   2. At the LOADER prompt, enter *bye* to reinitialize the PCIe cards and other components and let the controller reboot.

Step 6: Return the failed part to NetApp

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

Replace system management module - AFF A1K

Replace the System Management module in your AFF A1K system when it becomes defective or its firmware is corrupted. The replacement process involves shutting down the controller, replacing the failed System Management module, rebooting the controller, updating the license keys, and returning the failed part to NetApp.

The System Management module, located at the back of the controller in slot 8, contains onboard components for system management, as well as ports for external management. The target controller must be shut down to replace an impaired System Management module or replace the boot media.

The System Management module has the following onboard components:

- Boot media, allowing boot media replacement without removing the controller module.
- BMC
- Management switch

The System Management module also contains the following ports for external management:

- RJ45 Serial
- USB Serial (Type-C)
- USB Type-A (Boot recovery)
- e0M RJ45 Ethernet

Before you begin

- Make sure all other system components are working properly.
- Make sure that the partner controller is able to take over the impaired controller.
- Make sure you replace the failed component with a replacement component you received from NetApp.

About this task

This procedure uses the following terminology:

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Replace the impaired System Management module

Replace the impaired system management module.

Steps

1. Remove the System Management module:

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

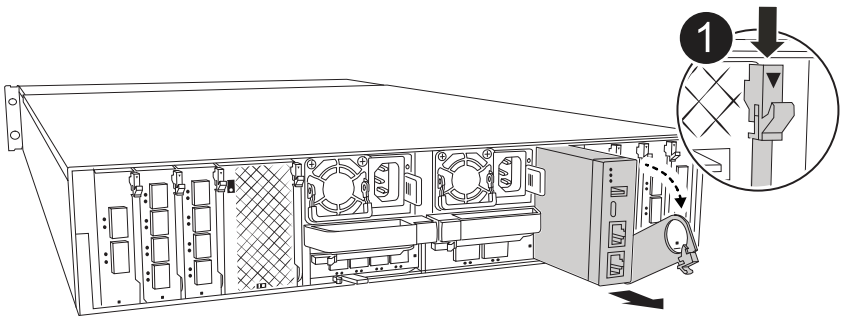

1	System Management module cam latch

   a. If you are not already grounded, properly ground yourself.
   b. Unplug the power supply cables from the PSUs.
2. Remove the System Manage module
   a. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Disconnect the power cords from the PSU for the impaired controller.
   c. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   d. Depress the cam button on the System Management module.
   e. Rotate the cam lever down as far as it will go.

f. Loop your finger into the hole on the cam lever and pull the module straight out of the system.

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

3. Move the boot media to the replacement System Management module:

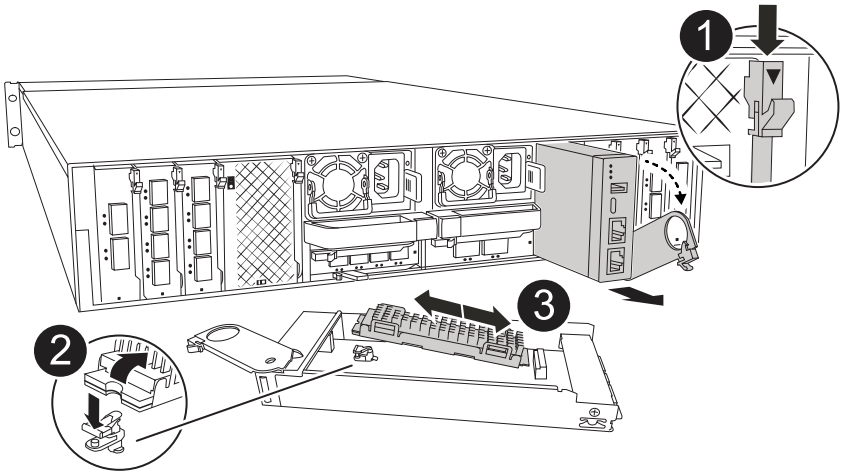

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue boot media locking button in the impaired System Management module.

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

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

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

b. Rotate the boot media down until it touches the locking button.

c. Depress the blue locking and rotate the boot media all the way down and release the blue locking button.

5. Install the replacement System Management module into the enclosure:

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

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

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

7. Recable the System Management module.

Step 3: Reboot the controller module

Reboot the controller module.

Steps

1. Plug the power cables back into the PSU.

   The system will begin to reboot, typically to the LOADER prompt.

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

Step 4: Install licenses and register serial number

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

About this task

Until you install license keys, features requiring standard licenses continue to be available to the node. However, if the node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the for the node as soon as possible.

Before you begin

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

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

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

Steps

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

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

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

4. Register the system serial number with NetApp Support.
    - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
    - If AutoSupport is not enabled, call NetApp Support to register the serial number.

Step 5: Return the failed part to NetApp

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

AFF A70 and AFF A90 systems

Install and setup

Installation and configuration workflow - AFF A70 and AFF A90

To install and configure your AFF A70 or AFF A90 system, you review the hardware requirements, prepare your site, install and cable the hardware components, power on the system, and set up your ONTAP cluster.

Review installation requirements

Review the equipment and tools needed to install your storage system and storage shelves and review the lifting and safety precautions.

Prepare to install the AFF A70 or AFF A90 storage system

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

Install the hardware for the AFF A70 or AFF A90 storage system

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system in the cabinet or telco rack. Next, slide the shelves onto the rails. Finally, attach cable management devices to the rear of the storage system for organized cable routing.

Cable the controllers and storage shelves for AFF A70 or AFF A90 storage system

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

Power on the AFF A70 or AFF A90 storage system

Before you power on the controllers, power on each NS224 shelf and assign a unique shelf ID to ensure each shelf is uniquely identified within the setup, connect the laptop or console to the controller, and then connect the controllers to the power sources.

6 Set up your cluster

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

Installation requirements - AFF A70 and AFF A90

Review the equipment needed and the lifting precautions for your AFF A70 or AFF A90 storage system and storage shelves.

Equipment needed for install

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

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

Lifting precautions

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

Storage system weight

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

An A1K storage system can weigh up to 62.83 lbs (28.5 kg). To lift the storage system, use two people or a hydraulic lift.

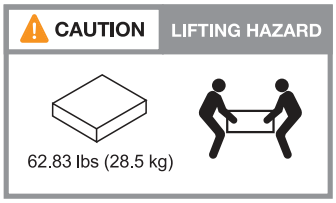

The storage system can weigh up to 151.68 lbs (68.8 kg). To lift the storage system, use four people or a hydraulic lift.

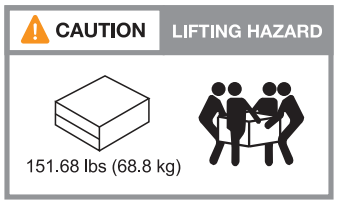

Shelf weight

Take the necessary precautions when moving or lifting your shelf.

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep

all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

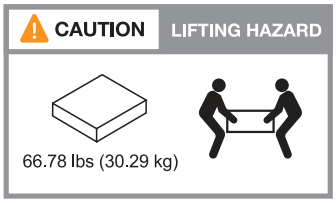

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

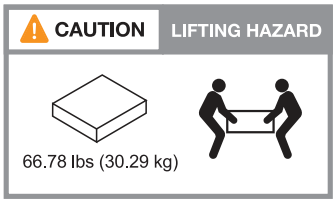

Related information

- Safety information and regulatory notices

What's next?

After you've reviewed the hardware requirements, you prepare to install your AFF A70 or AFF A90 storage system.

Prepare to install - AFF A70 and AFF A90

Prepare to install your AFF A70 or AFF A90 storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the system to access support benefits.

Step 1: Prepare the site

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

Steps

1. Use NetApp Hardware Universe to confirm that your site meets the environmental and electrical requirements for your storage system.
2. Make sure you have adequate cabinet or rack space for your storage system, shelves, and any switches:
   - 4U in an HA configuration
   - 2U for each NS224 storage shelf
3. Install any required network switches.

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

Step 2: Unpack the boxes

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

Steps

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

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

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

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

Hardware

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

Cables

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

Step 3: Register your storage system

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

Steps

1. Locate the System Serial Numbers (SSN) for every controller being installed. You can find the serial numbers in the following locations:
2. You can find the serial numbers in the following locations:
   - On the packing slip
   - In your confirmation email
   - On each controller's System Management module

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

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

What's next?

After you've prepared to install your AFF A70 or AFF A90 hardware, you install the hardware for your AFF A70 or AFF A90 storage system.

Install the hardware - AFF A70 and AFF A90

After you prepare to install your AFF A70 or AFF A90 storage system, install the hardware for the system. First, install the rail kits. Then install and secure your platform in a cabinet or telco rack.

Skip this step if your cabinet is pre-populated.

Before you begin

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

Steps

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

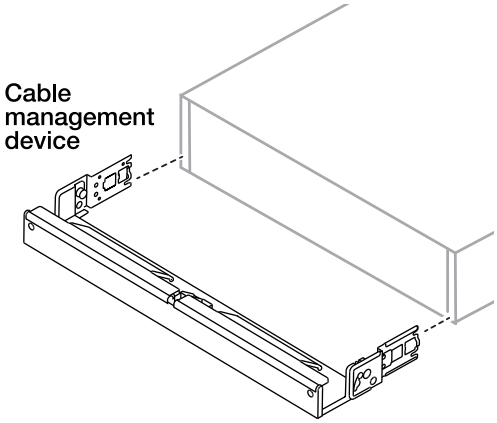

5. Install and secure the shelf as needed.
   a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

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

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

      What's next?

      After you've installed the hardware for your AFF A70 or AFF A90 system, you cable the hardware for your AFF A70 or AFF A90 storage system.

Cable the hardware - AFF A70 and AFF A90

After you install the rack hardware for your AFF A70 or AFF A90 storage system, install the network cables for the controllers, and connect the cables between the controllers and storage shelves.

Before you begin

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

About this task

- These procedures show common configurations. The specific cabling depends on the components ordered for your storage system. For comprehensive configuration and slot priority details, see NetApp Hardware Universe.
- The I/O slots on AFF A70 and AFF A90 controllers are numbered 1 through 11.

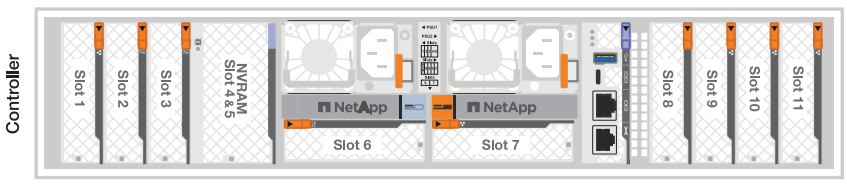

- The cabling graphics have arrow icons showing the proper orientation (up or down) of the cable connector pull-tab when inserting a connector into a port.

  As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it

over and try again.

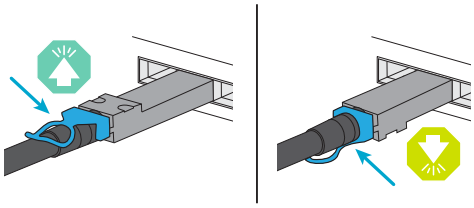

- If cabling to an optical switch, insert the optical transceiver into the controller port before cabling to the switch port.

Step 1: Connect the storage controllers to your network

Cable the controllers to your ONTAP cluster. This procedure differs depending on your storage system model and I/O module configuration.

The cluster interconnect traffic and the HA traffic share the same physical ports.

Switchless cluster cabling

Use the the Cluster/HA interconnect cable to connect to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A to port e1a on Controller B.
2. Connect port e7a on Controller A to port e7a on Controller B.

Cluster/HA interconnect cables

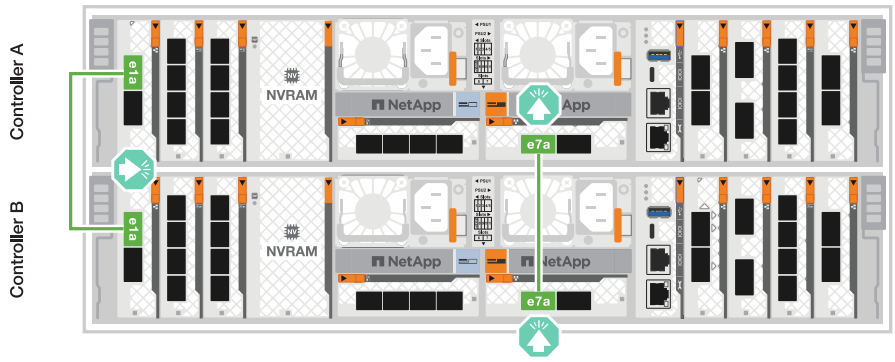

Switched cluster cabling

Use the 100 GbE cable to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A and port e1a on Controller B to cluster network switch A.
2. Connect port e7a on Controller A and port e7a on Controller B to cluster network switch B.

100 GbE cable

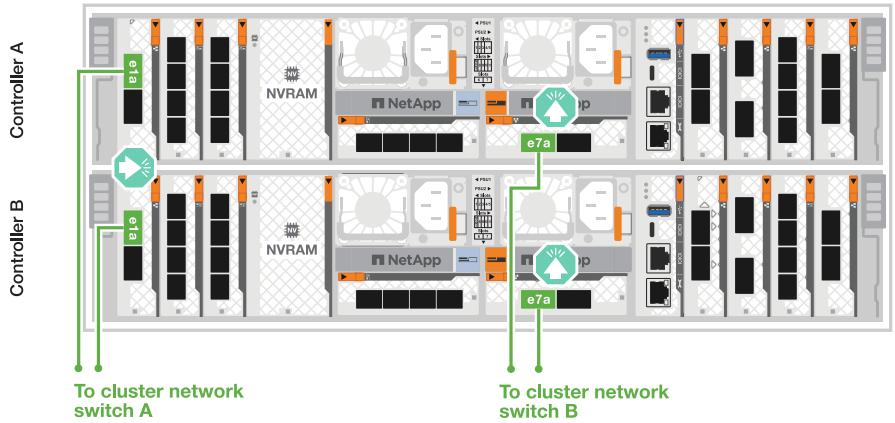

Step 2: Cable the host network connections

Connect the Ethernet module ports to your host network.

The following are some typical host network cabling examples. See NetApp Hardware Universe for your specific system configuration.

Steps

1. Connect ports e9a and e9b to your Ethernet data network switch.

   For maximum system performance for cluster and HA traffic, do not use ports e1b and e7b ports for host network connections. Use a separate host card to maximize performance.

   100 GbE cable

   

   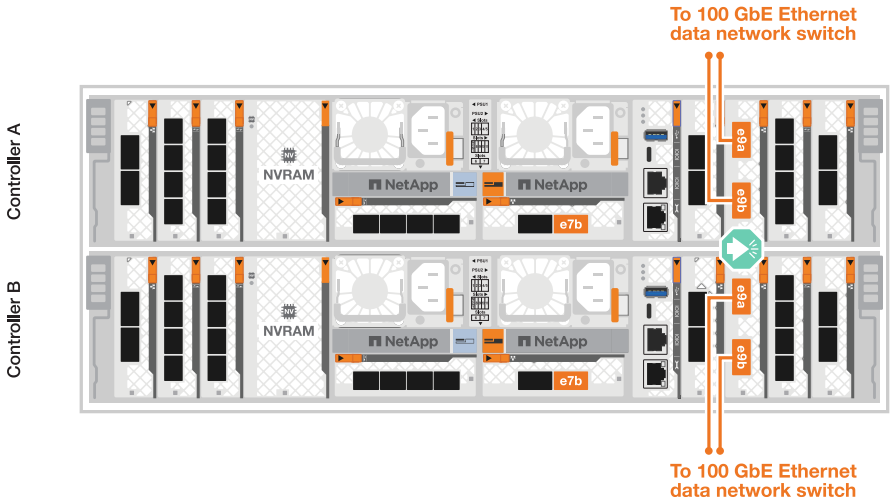

2. Connect your 10/25 GbE host network switches.

   4-ports, 10/25 GbE Host

   

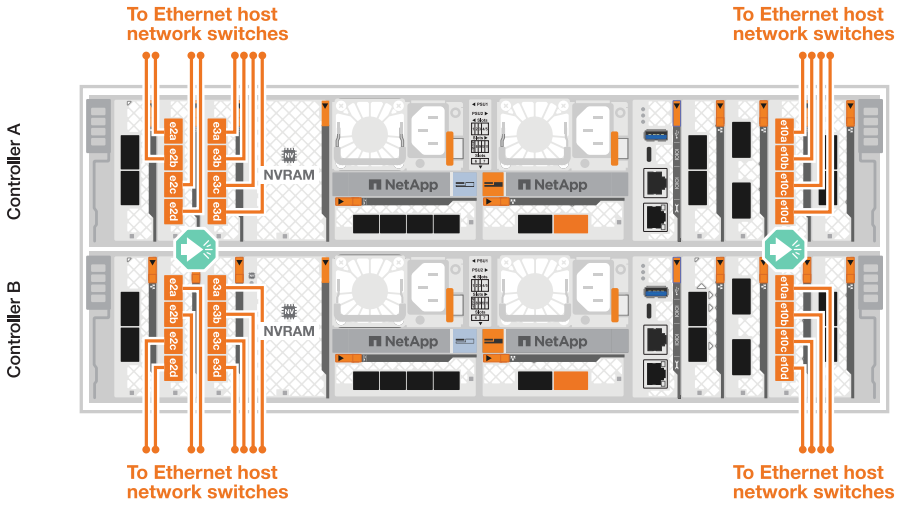

Step 3: Cable the management network connections

Use the 1000BASE-T RJ-45 cables to connect the management (wrench) ports on each controller to the management network switches.

1000BASE-T RJ-45 cables

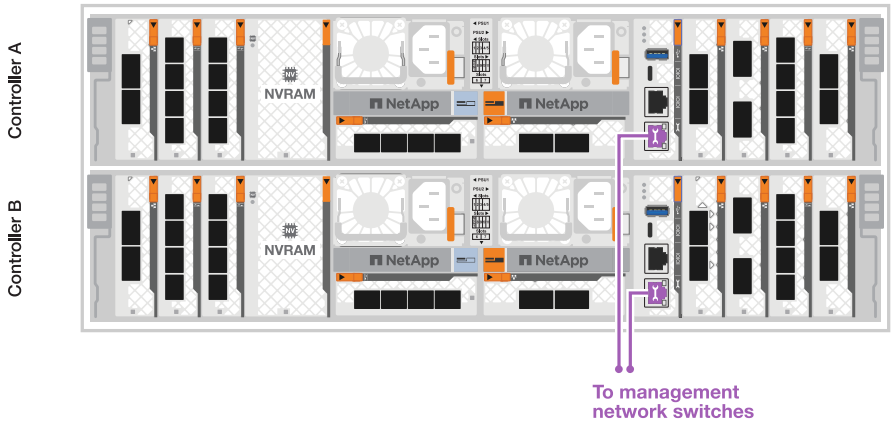

Do not plug in the power cords yet.

Step 4: Cable the shelf connections

The following cabling procedures show how to connect your controllers to a storage shelf. Choose one of the following cabling options that matches your setup.

For the maximum number of shelves supported for your storage system and for all of your cabling options, see NetApp Hardware Universe.

About this task

The AFF A70 and 90 storage systems support NS224 shelves with either the NSM100 or NSM100B module. The major differences between the modules are:

- NSM100 shelf modules use built-in ports e0a and e0b.

- NSM100B shelf modules use ports e1a and e1b in slot 1.

The following cabling example shows NSM100 modules in the NS224 shelves when referring to shelf module ports.

Option 1: One NS224 storage shelf

Connect each controller to the NSM modules on the NS224 shelf. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables

Steps

1. Connect controller A port e11a to NSM A port e0a.
2. Connect controller A port e11b to port NSM B port e0b.

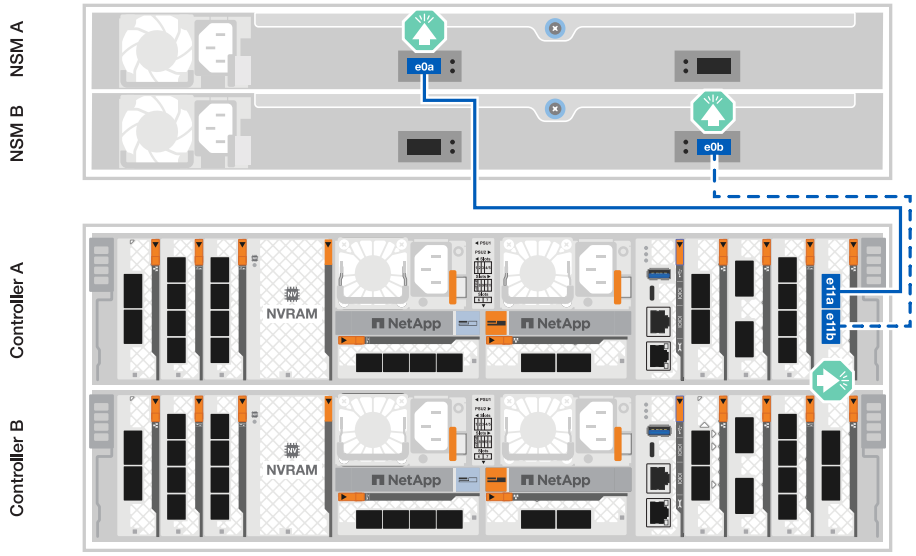

3. Connect controller B port e11a to NSM B port e0a.
4. Connect controller B port e11b to NSM A port e0b.

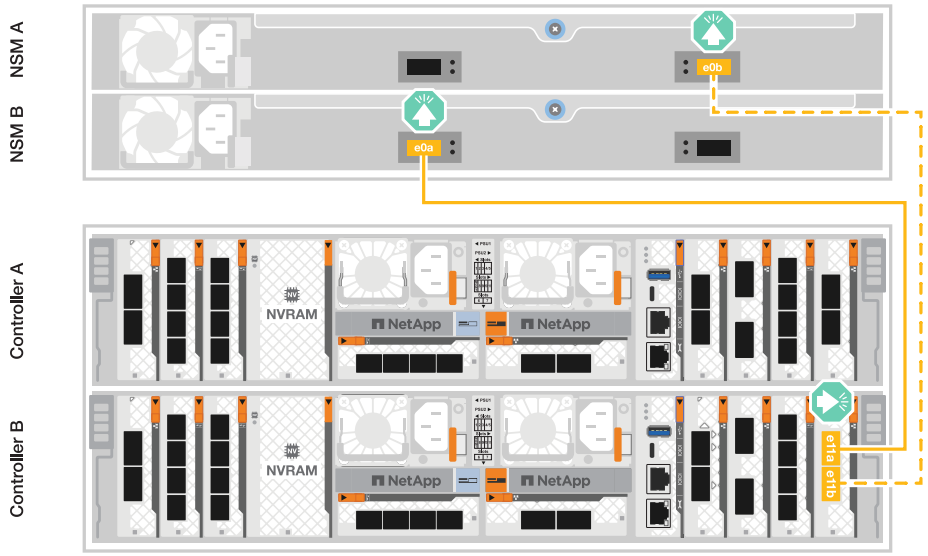

Option 2: Two NS224 storage shelves

Connect each controller to the NSM modules on both NS224 shelves. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables

Steps

1. On controller A, connect the following ports:
   a. Connect port e11a to shelf 1, NSM A port e0a.
   b. Connect port e11b to shelf 2, NSM B port e0b.
   c. Connect port e8a to shelf 2, NSM A port e0a.
   d. Connect port e8b to shelf 1, NSM B port e0b.

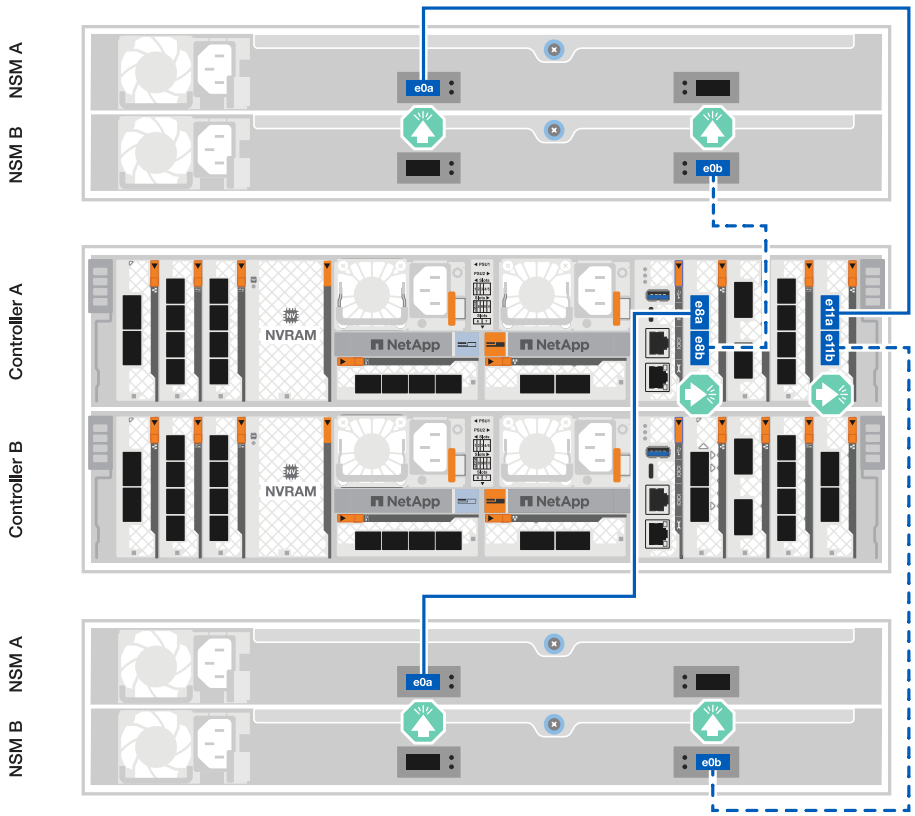

2. On controller B, connect the following ports:
   a. Connect port e11a to shelf 1, NSM B port e0a.
   b. Connect port e11b to shelf 2, NSM A port e0b.
   c. Connect port e8a to shelf 2, NSM B port e0a.
   d. Connect port e8b to shelf 1, NSM A port e0b.

What's next?

After you've cabled the hardware for your AFF A70 or AFF A90 system, you power on the AFF A70 or AFF A90 storage system.

Power on the storage system - AFF A70 and AFF A90

After you install the rack hardware for your AFF A70 or AFF A90 storage system and install the cables for the controllers and storage shelves, you should power on your storage shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

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

Before you begin

Make sure you have a paperclip or narrow tipped ball point pen for setting NS224 storage shelf IDs.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed shelf ID of 00.

- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, and then plug them back in) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, and then connecting the power cords to power sources on different circuits.

   The shelf powers on and boots automatically when plugged into the power source.

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

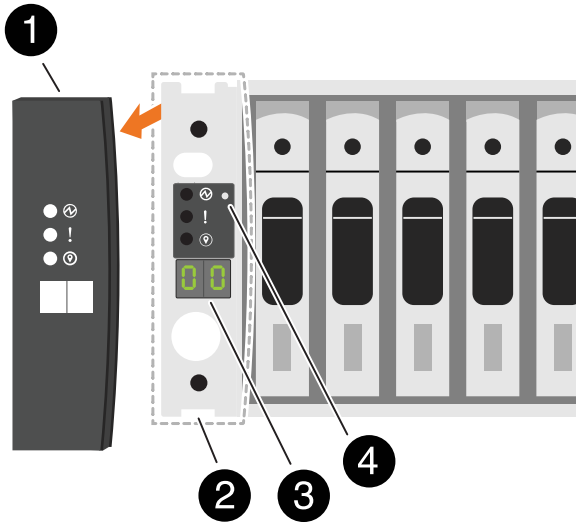

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID number
4	Shelf ID button

3. Change the first number of the shelf ID:
   a. Insert the straightened end of a paperclip or narrow tipped ball point pen into the small hole to press the shelf ID button.
   b. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

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

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

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

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

      The first number continues to blink.

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

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

      The first number on the digital display stops blinking.

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

      The second number continues to blink.

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

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

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

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.
   a. Unplug the power cord from both power supplies on the shelf.
   b. Wait 10 seconds.
   c. Plug the power cords back into the shelf power supplies to complete the power cycle.

      A power supply is powered on as soon as the power cord is plugged in. Its bicolored LED should illuminate green.

7. Replace the left end cap.

Step 2: Power on the controllers

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

Steps

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

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

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

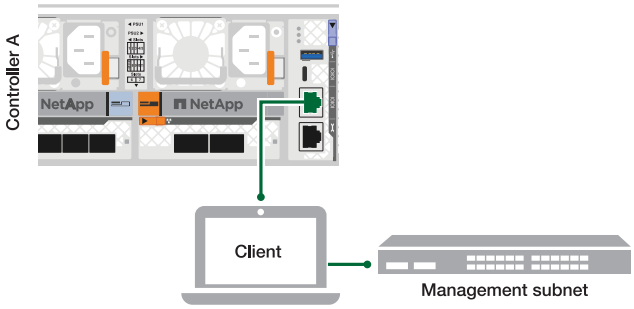

2. Assign a TCP/IP address to the laptop, using one that is on the management subnet.
3. Plug the two power cords into the controller power supplies, and then connect them to power sources on different circuits.

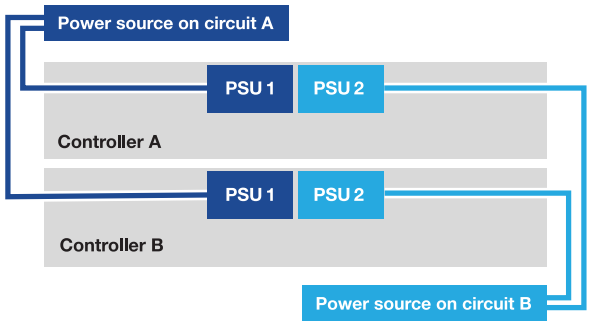

   - The system begins to boot. Initial booting might take up to eight minutes.
   - The LEDs flash on and the fans start, which indicates that the controllers are powering on.
   - The fans might be very noisy when they first start up. The fan noise during start-up is normal.
4. Secure the power cords using the securing device on each power supply.
   :a1k-a70-90!:

What's next?

After you've turned on your AFF A70 or AFF A90 storage system, you set up your cluster.

Maintain

Overview of the maintenance procedures - AFF A70 and AFF A90

Maintain the hardware of your AFF A70 and AFF A90 storage systems to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the AFF A70 and AFF A90 systems have already been deployed as a storage node in the ONTAP environment.

System components

For the AFF A70 and AFF A90 storage systems, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media- manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot the image from a USB drive and restore the configuration from the partner node.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller.
NVRAM	The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module.
NV battery	The NV battery is responsible for providing power to the NVRAM module while data in-flight is being destaged to flash memory after a power loss.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.
System Management module	The System Management module provides the interface between the controller and a console or laptop for controller or system maintenance purposes. The System management module contains the boot media and stores the system serial number (SSN).

Boot media automated recovery workflow - AFF A70 and AFF A90

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on the partner node to reinstall ONTAP on the replacement boot media in your AFF A70 or AFF A90 storage system.

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

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

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

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

3. **Replace the boot media**

Remove the failed boot media from the System Management module and install the replacement boot media.

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

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

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

Requirements for automated boot media recovery - AFF A70 and AFF A90

Before replacing the boot media in your AFF A70 or AFF A90 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming the cluster ports on the impaired controller are working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

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

Review the following requirements.

- You must replace the failed boot media with a replacement boot media you received from NetApp.
- The cluster ports are used to communicate between the two controllers during the automated boot recovery process. Make sure that the cluster ports on the impaired controller are working properly.
- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg
  - /cfcard/kmip/certs/client.crt
  - /cfcard/kmip/certs/client.key
  - /cfcard/kmip/certs/CA.pem
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

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

Shut down the controller for automated boot media recovery - AFF A70 and AFF A90

Shut down the impaired controller in your AFF A70 or AFF A90 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

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

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

About this task

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

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

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

Steps

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

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

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

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

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

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

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

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

What's next

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

Replace the boot media for automated boot recovery - AFF A70 and AFF A90

The boot media in your AFF A70 or AFF A90 storage system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media in the System Management module, and then reinstalling the System Management module.

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

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

Steps

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

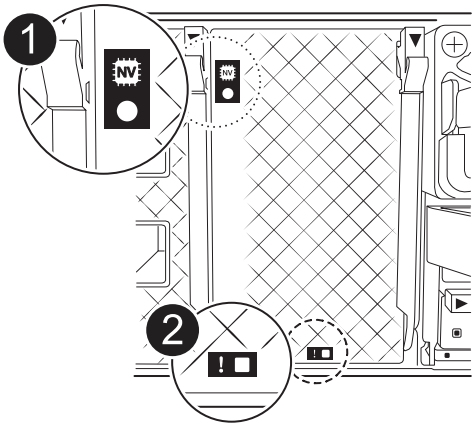

1	NVRAM status LED
2	NVRAM attention LED

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

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller's PSUs.

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

   a. Remove any cables connected to the System Management module. Make sure to label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   c. Depress the system management cam button.
   The cam lever moves away from the chassis.
   d. Rotate the cam lever all the way down and remove the System Management module from the controller module.
   e. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:

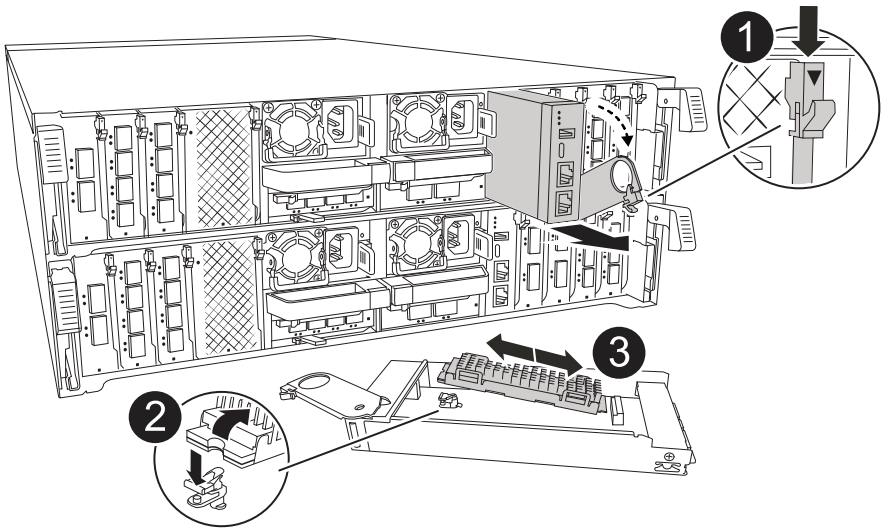

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue locking button.

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

5. Install the replacement boot media into the System Management module:

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

b. Rotate the boot media down toward the locking button.

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

6. Reinstall the System Management module:

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

b. Recable the System Management module.

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

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

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

What's next

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

Automated boot media recovery from the partner node - AFF A70 and AFF A90

After installing the new boot media device in your AFF A70 or AFF A90 storage system, you can start the automated boot media recovery process to restore the configuration

from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

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

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

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

2. Monitor the boot media install recovery process.

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

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

   Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

   Show example of configuration error finding prompts

   ```
   Error when fetching key manager config from partner ${partner_ip}:
   ${status}

   Has key manager been configured on this system

   Is the key manager onboard
   ```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

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

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

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

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

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

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

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

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

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

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

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

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

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

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

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

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

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

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

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

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

5. If automatic giveback was disabled, reenable it:

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

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

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

What's next

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

Return the failed boot media part to NetApp - AFF A70 and AFF A90

If a component in your AFF A70 or AFF A90 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF A70 and AFF A90

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the AFF A70 or AFF A90 system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

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

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

Review requirements to replace the boot media

Review the requirements for replacing the boot media.

Check encryption key support and status

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

Shut down the controller

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

Replace the boot media

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

Boot the recovery image

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

6 Restore encryption

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

7 Return the failed part to NetApp

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

Requirements for manual boot media recovery - AFF A70 and AFF A90

Before replacing the boot media in your AFF A70 or AFF A90 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

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

File preparation

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

Component replacement

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

Controller identification

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

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

What's next?

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

Check encryption support for manual boot media recovery - AFF A70 and AFF A90

To ensure data security on your AFF A70 or AFF A90 storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

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

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

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

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

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

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

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

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

What's next?

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

Shut down the controller for manual boot media recovery - AFF A70 and AFF A90

Shut down the impaired controller in your AFF A70 or AFF A90 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

What's next?

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

Replace the boot media and prepare for manual boot recovery - AFF A70 and AFF A90

The boot media in your AFF A70 or AFF A90 system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

Step 1: Replace the boot media

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

Steps

1. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
2. Unplug the controller's PSUs.

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

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

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

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

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

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

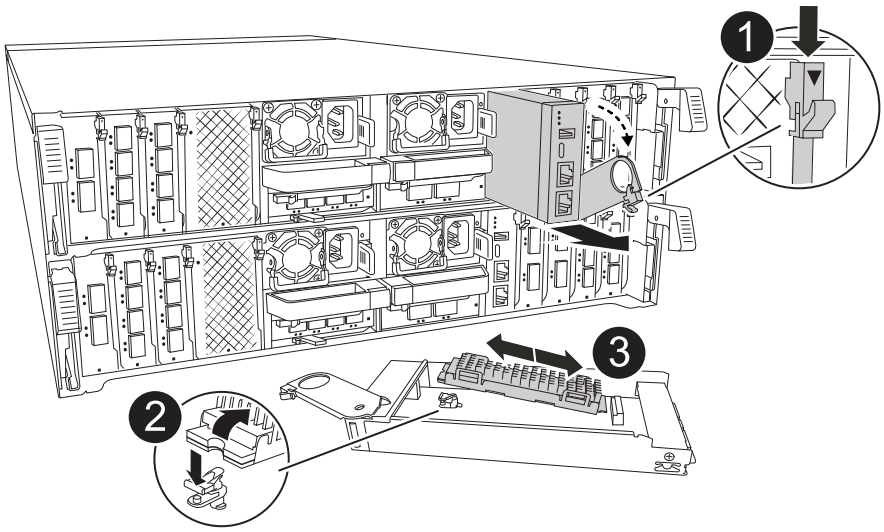

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue locking button.

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

4. Install the replacement boot media into the System Management module:

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

b. Rotate the boot media down toward the locking button.

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

5. Reinstall the System Management module:

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

b. Recable the System Management module.

Step 2: Transfer the boot image to the boot media

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

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- Download a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site. Use the `version -v` command to display if your version of ONTAP supports NVE. If the command output displays `<10no- DARE>`, your version of ONTAP does not support NVE.
  - If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption,

as indicated in the download button.

- If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

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

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

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

   c. Remove the USB flash drive from your laptop.
2. Insert the USB flash drive into the USB-A port on the System Management module.

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

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

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

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

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

What's next?

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

Manual boot media recovery from a USB drive - AFF A70 and AFF A90

After installing the new boot media device in your AFF A70 or AFF A90 system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Steps

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

   The image is downloaded from the USB flash drive.

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

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

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

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

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

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

What's next?

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

Restore encryption keys after manual boot recovery - AFF A70 and AFF A90

Restore encryption on the replacement boot media in your AFF A70 or AFF A90 system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

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

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

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

3. Confirm that you want to continue the recovery process.

   Show example prompt

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

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

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

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

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

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

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

Successfully recovered keymanager secrets.

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

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

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

Show example prompt

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


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

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

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

8. From the partner node, giveback the partner controller by entering the following command.

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

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

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

Show example prompt

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

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

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

11. Ensure that all keys are synced by entering the following command.

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

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

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

13. Restore automatic giveback, if you disabled it, by entering the following command.

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

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

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

Option 2: Restore the External Key Manager configuration

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

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

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

   Show example boot menu

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

3. When prompted, confirm you have gathered the required information.

   Show example prompt

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

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

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

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

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

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

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

After you enter the client and server information, the recovery process completes.

Show example

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

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

Show example prompt

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

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

6. Restore automatic giveback if you disabled it.

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

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

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

What's next?

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

Return the failed part to NetApp - AFF A70 and AFF A90

If a component in your AFF A70 or AFF A90 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - AFF A70 and AFF A90

Get started with replacing the chassis of your AFF A70 or AFF A90 storage system by reviewing the replacement requirements, shutting down the controllers, replacing the chassis, and verifying system operations.

1 Review the chassis replace requirements

Review the chassis replacement requirements.

2 Prepare for chassis replace

Prepare to replace the chassis by locating the system, gathering system credentials and necessary tools, verifying the replacement chassis was received, and labeling the system cables.

3 Shut down the controllers

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

4 Replace the chassis

Replace the chassis by moving the components from the impaired chassis to the replacement chassis.

5 Complete the chassis replacement

Complete the chassis replacement by bringing the controllers up, giving back the controllers, and returning the failed chassis to NetApp.

Requirements to replace the chassis - AFF A70 and AFF A90

Before replacing the chassis in your AFF A70 or AFF A90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have local administrator credentials for ONTAP, the correct replacement chassis, and the necessary tools.

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

Review the following requirements.

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

What’s next?

After you've reviewed the requirements to replace the chassis, you need to prepare to replace the chassis.

Prepare to replace the chassis - AFF A70 and AFF A90

Prepare to replace the impaired chassis in your AFF A70 or AFF A90 system by identifying the impaired chassis, verifying the replacement components, and labeling the cables and controller modules.

Step 1: Locate and monitor your system

You should open a console session and save sessions logs for future reference, and also turn on the system location LED to find the impaired chassis.

Steps

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

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

      The Location LED remains lit for 30 minutes.

Step 2: Verify replacement components

You should verify that you received the necessary components, remove them from packaging, and save the packaging.

Steps

1. Before opening the packaging, you should look at the packaging label and verify:
   - Component part number.
   - Part description.
   - Quantity in the box.
2. Remove the contents from the packaging and use the packaging to returning the failed component to NetApp.

Step 3: Label the cables and controller modules

You should label the cables and controller modules before removing them from the controller modules or chassis.

Steps

1. Label all the cables associated with the storage system. This aids recabling later in this procedure.
2. Label the controller modules.
3. If you are not already properly grounded, ground yourself.

What's next?

After you've prepared to replace your AFF A70 or AFF A90 chassis hardware, you need to shut down the controllers.

Shut down the controllers to replace the chassis - AFF A70 and AFF A90

Shut down the controllers in your AFF A70 or AFF A90 storage system to prevent data loss and ensure system stability when replacing the chassis.

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

Before you begin

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

Steps

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

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

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

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

6. Exit the cluster shell:

   ```
   exit
   ```

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

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

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

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

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

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

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

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

What's next?

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

Replace the chassis - AFF A70 and AFF A90

Replace the chassis of your AFF A70 or AFF A90 system when a hardware failure requires it. The replacement process involves removing the controllers and power supply units (PSUs), removing the drives, installing the replacement chassis, and reinstalling the chassis components.

Step 1: Remove the PSUs and cables

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

Steps

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

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

   c. Remove the PSU from the controller by rotating the PSU handle up so that you can pull the PSU out, press the PSU locking tab, and then pull PSU out of the controller module.

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

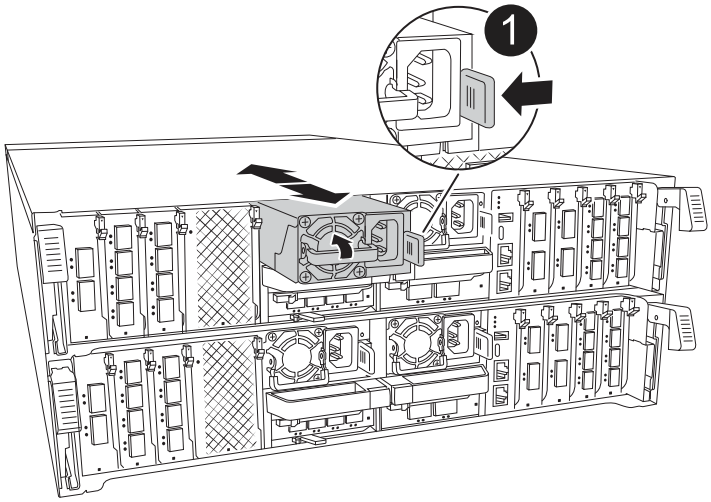

1	Terracotta PSU locking tab

d. Repeat these steps for the remaining PSUs.

2. Remove the cables:

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

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

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

Step 2: Remove the controller modules and drives

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

Steps

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

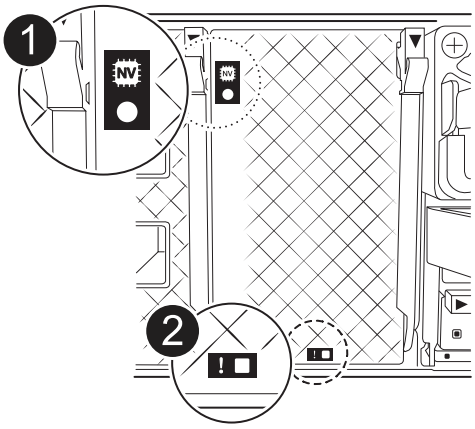

1	NVRAM status LED

2	NVRAM attention LED

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

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

      The controller module moves slightly out of the chassis.

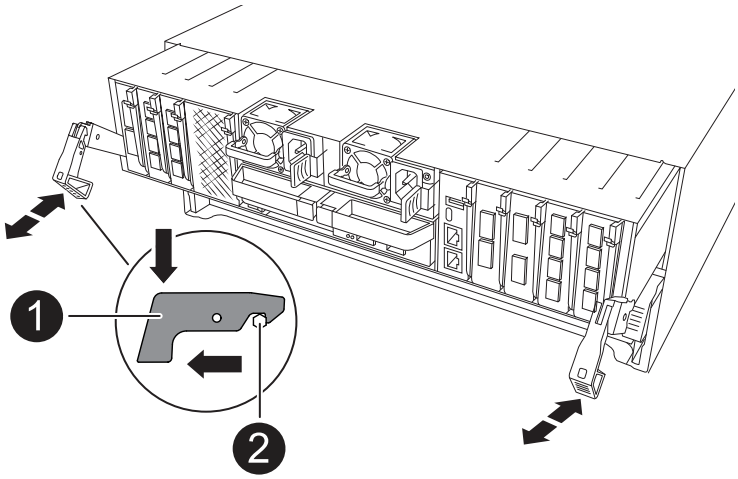

1	Locking latch
2	Locking pin

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

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

   c. Repeat these steps for the second controller module.
3. Remove the drives:
   a. Gently remove the bezel from the front of the system.
   b. Press the release button at the top of the drive carrier face below the LEDs.
   c. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

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

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

d. Keep track of which drive bay each drive was in and set the drive aside on a static-free cart or table.

e. Repeat this step for the remaining drives in the chassis.

Step 3: Replace the impaired chassis

Remove the impaired chassis and install the replacement chassis.

Steps

1. Remove the impaired chassis:
   a. Remove the screws from the chassis mount points.
   b. Using two people or a lift, slide the impaired chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
2. Install the replacement chassis:
   a. Using two people or a lift, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
   b. Slide the chassis all the way into the equipment rack or system cabinet.
   c. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.

Step 4: Install the chassis components

After the replacement chassis is installed, you need to install the controller modules, recable them, and then reinstall the drives and PSUs.

Steps

1. Beginning with the bottom controller module, install the controller modules in the replacement chassis:
   a. Align the end of the controller module with the opening in the chassis, and then gently push the controller all the way into the chassis.
   b. Rotate the locking latches upward into the locked position.
   c. If you have not already done so, reinstall the cable management device and recable the controller.

      If you removed the media converters (QSFPs or SFPs), remember to reinstall them.

      Make sure that the cables are connected referencing the cable labels.
2. Reinstall the drives into their corresponding drive bays in the front of the chassis.
3. Install all four of the PSUs:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.
4. Reconnect the PSU power cables to all four of the PSUs.

a. Secure the power cable to the PSU using the power cable retainer.

   If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis and secure the power cable to the PSU with the thumbscrews.

The controller modules begin to boot as soon as PSUs are installed and power is restored.

What's next?

After you've replaced the impaired AFF A70 or AFF A90 chassis and reinstalled the components into it, you need to complete the chassis replacement.

Complete the chassis replacement - AFF A70 and AFF A90

Reboot the controllers, verify system health, and return the failed part to NetApp to complete the final step in the AFF A70 and AFF A90 chassis replacement procedure.

Step 1: Boot the controllers and give back the controllers

After the controllers reboot, boot ONTAP and give back the controllers.

Steps

1. Check the console output:
   a. If the controller boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.
   b. If the console displays `waiting for giveback` after the reboot, log into the partner controller and check that the replaced controller is ready for giveback with the `storage failover show` command.
2. Perform the giveback:
   a. Connect the console cable to the partner controller.
   b. Give back the controller with the `storage failover giveback -fromnode local` command.

Step 2: Verify storage system health

After the controller has given back the storage, you should check the overall health with Active IQ Config Advisor.

Steps

1. After the giveback is complete, run Active IQ Config Advisor to verify the health of the storage system.
2. Correct any issues you encounter.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - AFF A70 and AFF A90

Get started with replacing the controller in your AFF A70 or AFF A90 storage system by

shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

1. **Review the requirements to replace the controller**

To replace the controller module, you must meet certain requirements.

2. **Shut down the impaired controller**

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

3. **Replace the controller**

Replacing the controller includes removing the impaired controller, moving the FRU components to the replacement controller module, and then installing the replacement controller module in the enclosure.

4. **Restore and verify the system configuration**

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

5. **Recable and give back the controller**

Recable the controller and transfer the ownership of storage resources back to the replacement controller.

6. **Complete controller replacement**

Verify the LIfs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - AFF A70 and AFF A90

Before replacing the controller of your AFF A70 or AFF A90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements for replacing the controller module.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- Do not use this procedure for controller upgrades; instead, refer to the Choose your controller hardware upgrade procedure for guidance.
- If your system is in a MetroCluster configuration, you must review Choosing the correct recovery procedure

to determine whether you should use this procedure.

- You must replace the failed component with the field-replaceable unit (FRU) you received from NetApp.
- You must replace a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text log file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

What's next?

After you've the reviewed the requirements to replace your AFF A70 or AFF A90 controller, you need to shut down the impaired controller.

Shut down the impaired controller - AFF A70 and AFF A90

Shut down the controller in your AFF A70 or AFF A90 storage system to prevent data loss and ensure system stability when replacing the controller.

Shut down the controller module using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After you've shut down the controller, you need to replace the controller.

Replace the controller - AFF A70 and AFF A90

Replace the controller in your AFF A70 or AFF A90 system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

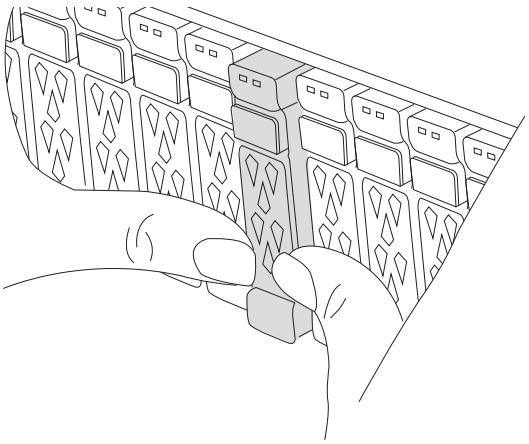

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

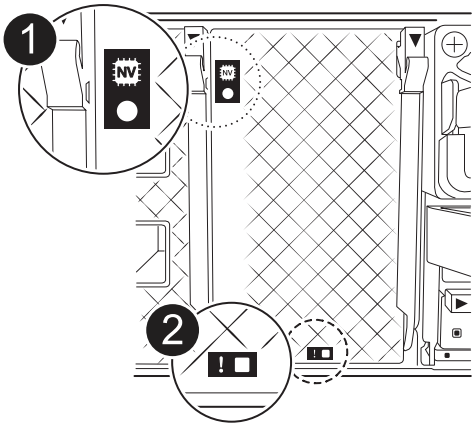

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

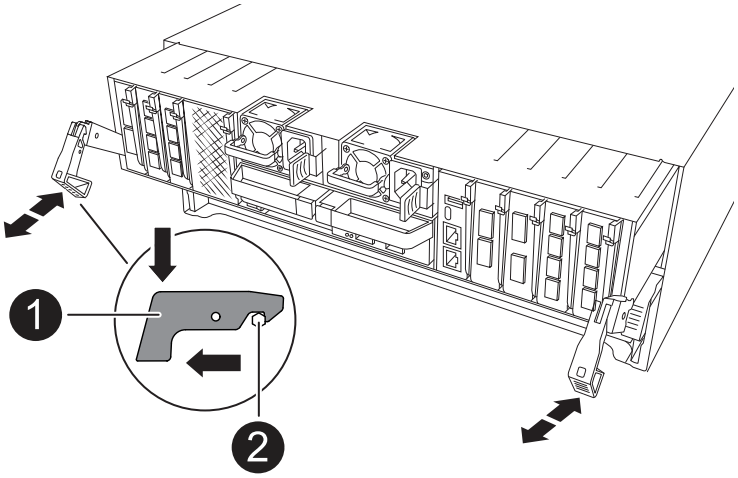

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 2: Move the power supplies

Move the power supplies to the replacement controller.

Steps

1. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

   The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

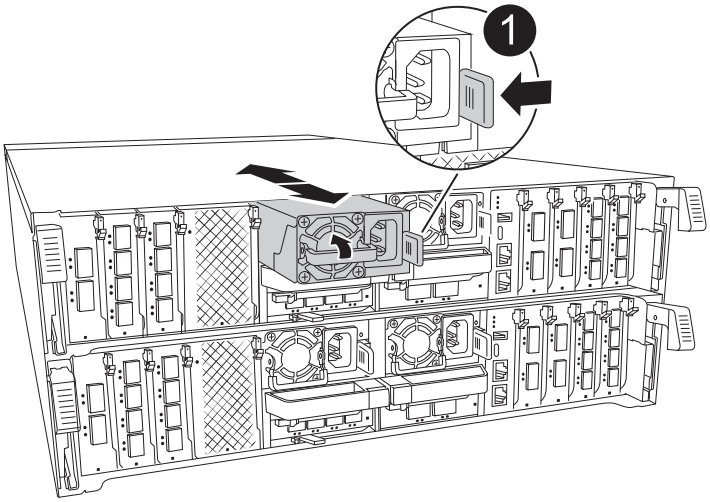

1	Terracotta PSU locking tab
2	Power supply

2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

> To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

Move the fans modules to the replacement controller module.

Steps

1. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

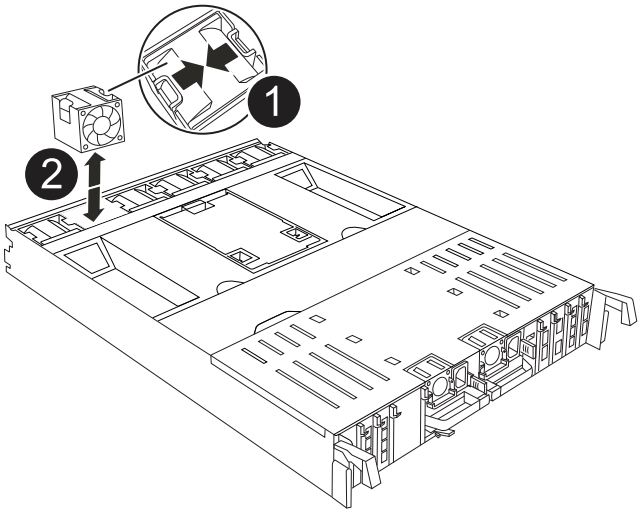

1	Fan locking tabs
2	Fan module

2. Move the fan module to the replacement controller module, and then install the fan module by aligning its edges with the opening in the controller module, and then sliding the fan module into the controller module until the locking latches click into place.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the NV battery

Move the NV battery to the replacement controller module.

Steps

1. Open the air duct cover in the middle of the controller module and locate the NV battery.

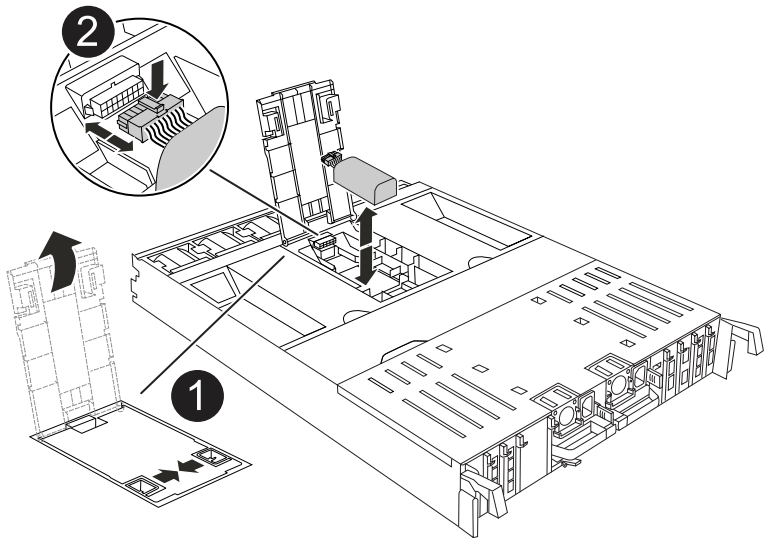

1	NV battery air duct
2	NV battery pack plug

Attention: The NV module LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module.
5. Move the battery pack to the replacement controller module and then install it in the replacement controller module:
   a. Open the NV battery air duct in the replacement controller module.
   b. Plug the battery plug into the socket and make sure that the plug locks into place.
   c. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   d. Close the NV battery air duct.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller module.

Steps

1. Open the controller air duct on the top of the controller.

a. Insert your fingers in the recesses at the far ends of the air duct.

b. Lift the air duct and rotate it upward as far as it will go.

2. Locate the system DIMMs on the motherboard, using the DIMM map on top of the air duct.

   The DIMM locations, by model, are listed in the following table:

Model	DIMM slot location
FAS70	3, 10, 19, 26
FAS90	3, 7, 10, 14, 19, 23, 26, 30

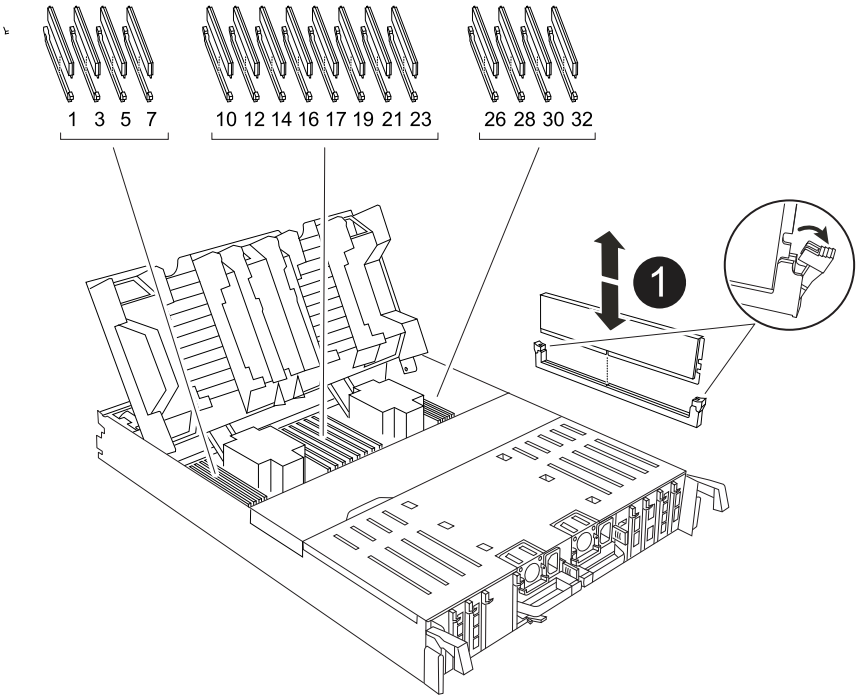

1	System DIMM

3. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.

4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Locate the slot on the replacement controller module where you are installing the DIMM.

6. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Repeat these steps for the remaining DIMMs.
9. Close the controller air duct.

Step 6: Move the I/O modules

Move the I/O modules to the replacement controller module.

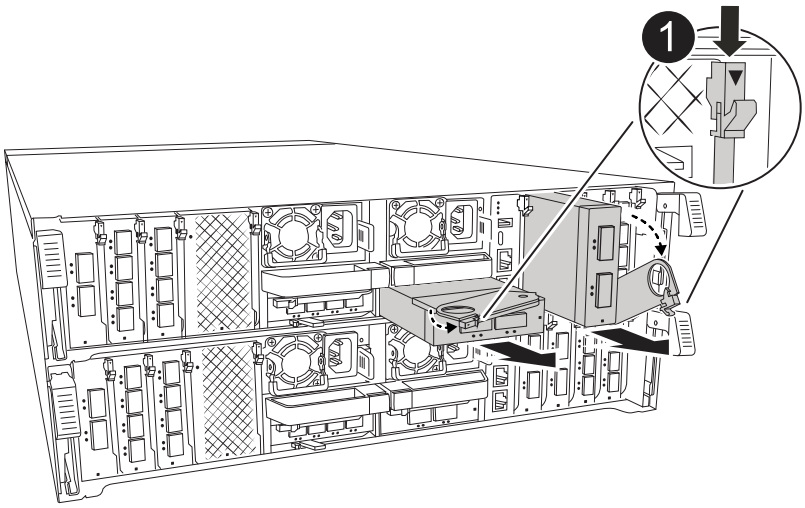

1	I/O module cam lever

Steps

1. Unplug any cabling on the target I/O module.

   Make sure that you label the cables so that you know where they came from.

2. Rotate the cable management arm down by pulling the buttons on the inside of the cable management arm and rotating it down.
3. Remove the I/O modules from the controller module:
   a. Depress the target I/O module cam latch button.
   b. Rotate the cam latch down as far as it will go. For horizontal modules, rotate the cam away from the module as far as it will go.
   c. Remove the module from the controller module by hooking your finger into the cam lever opening and pulling the module out of the controller module.

      Make sure that you keep track of which slot the I/O module was in.

   d. Install the replacement I/O module into the replacement controller module by gently sliding the I/O module into the slot until the I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.
4. Repeat these steps to move the remaining I/O modules, except the modules in slots 6 and 7, to the

replacement controller module.

> To move the I/O modules from slots 6 and 7, you must move the carrier containing these I/O modules from the impaired controller module to the replacement controller module.

5. Move the carrier containing the I/O modules in slots 6 and 7 to the replacement controller module:
   a. Press the button on the right-most handle on the carrier handle.
   ..Slide the carrier out of the impaired controller module insert it into the replacement controller module in the same position it was in the impaired controller module.
   b. Gently push the carrier all the way into the replacement controller module until it locks into place.

Step 7: Move the System Management module

Move the System Management module to the replacement controller module.

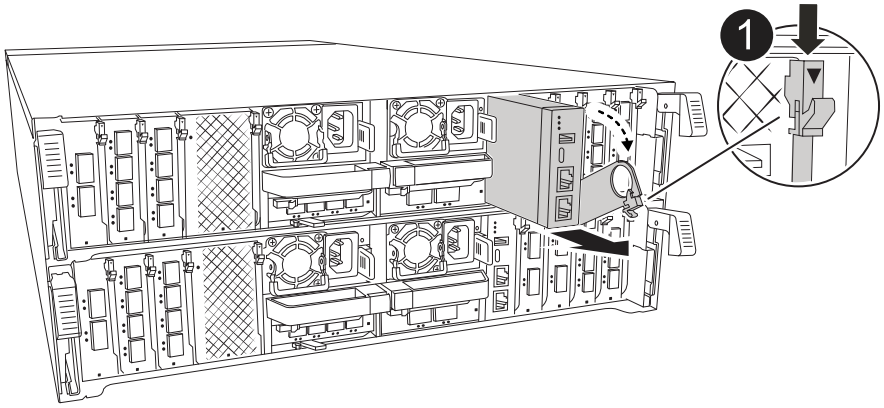

1	System Management module cam latch

Steps

1. Remove the System Management module from the impaired controller module:
   a. Depress the system management cam button.
   b. Rotate the cam lever all the way down.
   c. Loop your finger into the cam lever and pull the module straight out of the system.
2. Install the system management module into the replacement controller module in the same slot that it was in on the impaired controller module:
   a. Align the edges of the System Management module with the system opening and gently push it into the controller module.
   b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

Step 8: Move the NVRAM module

Move the NVRAM module to the replacement controller module.

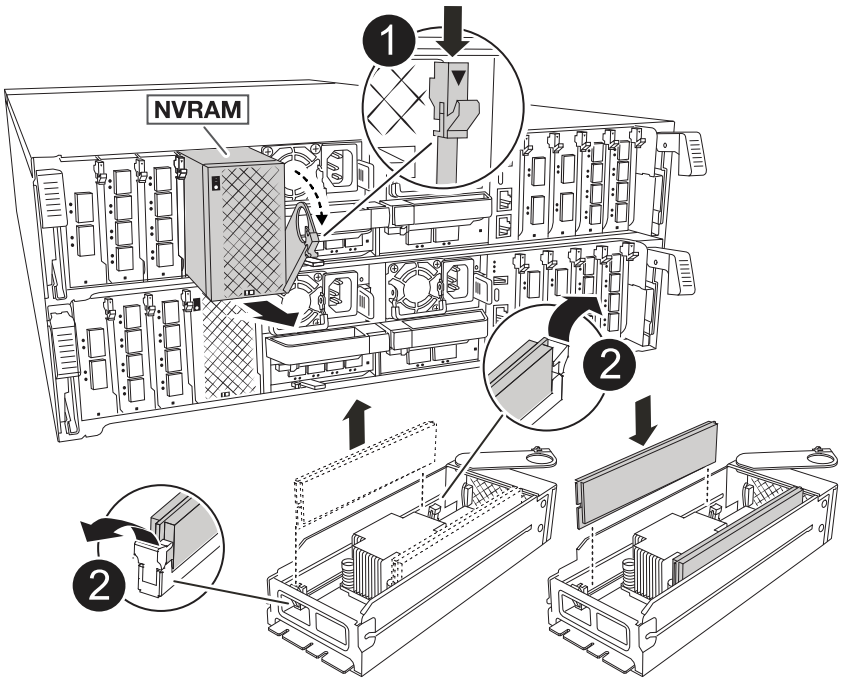

1	Cam locking button
2	DIMM locking tab

Steps

1. Remove the NVRAM module from the impaired controller module:
   a. Depress the cam latch button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.
   c. Remove the NVRAM module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
2. Install the NVRAM module into slot 4/5 in the replacement controller module:
   a. Align the module with the edges of the chassis opening in slot 4/5.
   b. Gently slide the module into the slot all the way, and then push the cam latch all the way up to lock the module in place.

Step 9: Install the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller

module halfway into the system.

> Do not completely insert the controller module in the chassis until instructed to do so.

3. Reinstall the cable management arm, if removed, but do not reconnect any cables to the replacement controller.
4. Plug the console cable into the console port of the replacement controller module and reconnect it to the laptop so that so that it receives console messages when it reboots.
5. Complete the reinstallation of the controller module:

   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   a. Rotate the locking latches upward into the locked position.
   b. Plug in the power supplies. The controller boots to the LOADER prompt as soon as power is restored.
6. From the LOADER prompt, enter `show date` to display the date and time on the replacement controller. Date and time are in GMT.

   > Time displayed is local time not always GMT and is displayed in 24hr mode.

7. Set the current time in GMT with the `set time hh:mm:ss` command. You can get the current GMT from the partner node the `date -u`command.
8. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

9. Plug the power cords into the power supplies.

   > If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

What's next?

After you've replaced the impaired AFF A70 or AFF A90 controller, you need to restore the system configuration.

Restore and verify the system configuration - AFF A70 and AFF A90

Verify that the controller's HA configuration is active and functioning correctly in your AFF A70 or AFF A90 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

Steps

1. Boot to maintenance mode: `boot_ontap maint`
   a. Enter `y` when you see *Continue with boot?*.

      If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

5. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha`

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed: `ha-config show`

Step 2: Verify disk list

You must verify the adapter list and paths to all your system disks.

Steps

1. Verify that the adapter lists the paths to all disks with the `storage show disk -p`.

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode: `halt`.

What's next?

After you've restored and verified the system configuration for your AFF A70 or AFF A90 system, you need to give back the controller.

Give back the controller - AFF A70 and AFF A90

Return control of storage resources to the replacement controller so your AFF A70 or AFF A90 system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`
   - If you encounter errors, contact NetApp Support.
9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

    c. Verify status of the keys: `security key-manager key query -restored false`

The output should show no results when when properly synchronized.

If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

After you've transferred the ownership of storage resources back to the replacement controller, you need to complete the controller replacement procedure.

Complete controller replacement - AFF A70 and AFF A90

To complete the controller replacement for your AFF A70 or AFF A90 system, first restore the NetApp Storage Encryption configuration (if necessary). Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, return the failed part to NetApp.

Step 1: Verify LIFs and check cluster health

Before returning the replacement node to service, verify that the logical interfaces are on their home ports, check the cluster health, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any logical interfaces are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A70 and AFF A90

Replace a DIMM in your AFF A70 or AFF A90 system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system

from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- Make sure all other components in the system are functioning properly; if not, you must contact technical support.
- Make sure you replace the failed component with a replacement component you received from NetApp.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

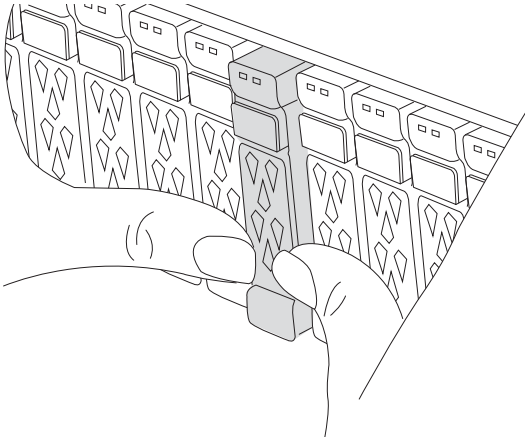

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

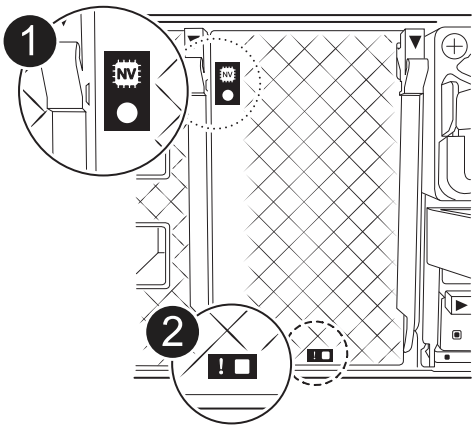

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

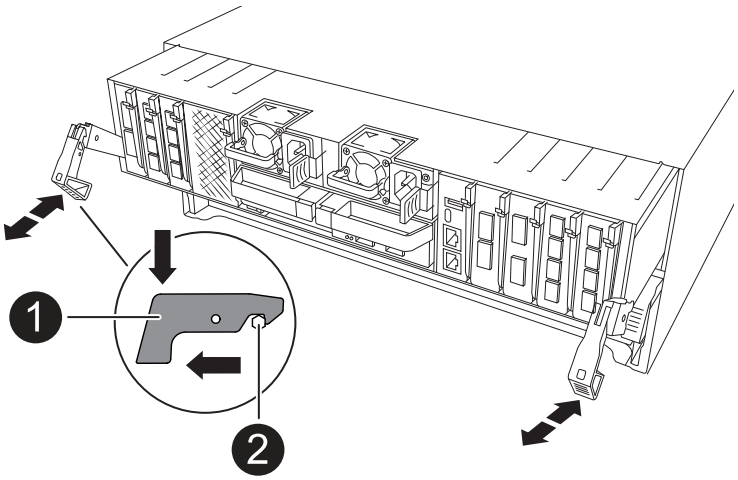

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the DIMM

To replace the DIMM, locate them inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
3. Locate the DIMMs on your controller module and identify the target DIMM.

   Use the FRU map on the controller airduct to locate the DIMM slot.

4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

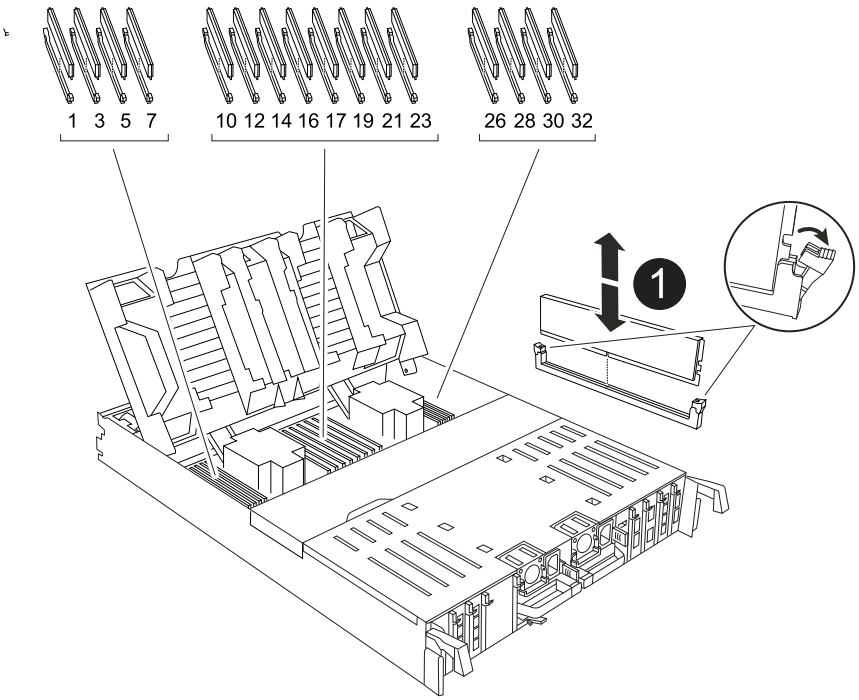

1	DIMM and DIMM ejector tabs

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the controller air duct.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

   Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward into the locked position.
5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

6. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode` *`impaired_node_name`*.

7. If automatic giveback was disabled, reenable it:

   `storage failover modify -node local -auto-giveback true`.

8. If AutoSupport is enabled, restore/unsuppress automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END.`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an SSD drive - AFF A70 and AFF A90

Replace a drive in your AFF A70 or AFF A90 system when a drive fails or requires an upgrade. The replacement process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed SSD drive nondisruptively while I/O is in progress.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed disk drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  It can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the disk depends on how the disk drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

When replacing several disk drives, you must wait 70 seconds between the removal of each failed disk drive and the insertion of the replacement disk drive to allow the storage system to recognize the existence of each new disk.

You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   > The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.
8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.

a. Display all unowned drives:

```
storage disk show -container-type unassigned
```

You can enter the command on either controller module.

b. Assign each drive:

```
storage disk assign -disk disk_name -owner node_name
```

You can enter the command on either controller module.

You can use the wildcard character to assign more than one drive at once.

c. Reenable automatic drive assignment if needed:

```
storage disk option modify -node node_name -autoassign on
```

You must reenable automatic drive assignment on both controller modules.

10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Replace a fan module - AFF A70 and AFF A90

Replace a fan module in your AFF A70 or AFF A90 system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

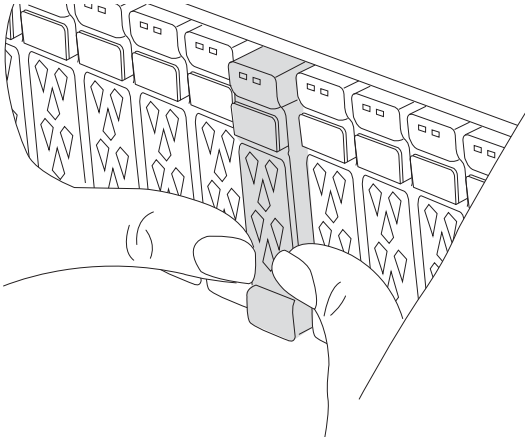

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

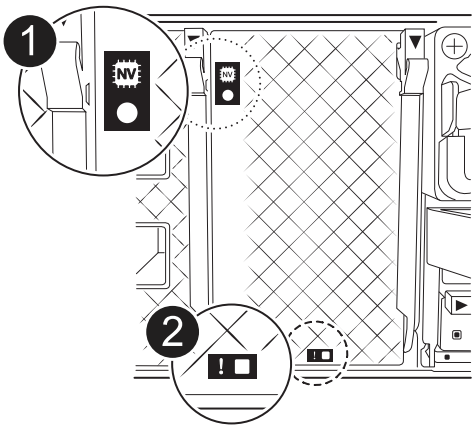

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

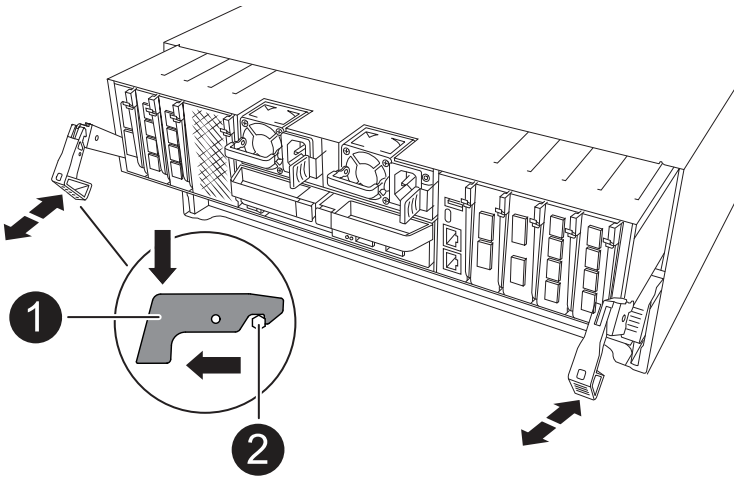

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

Steps

1. Identify the fan module that you must replace by checking the console error messages.
2. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

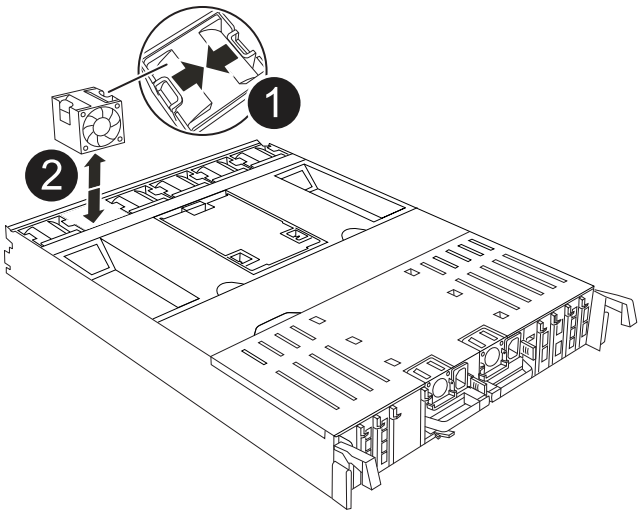

1	Fan locking tabs
2	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module until the locking latches click into place.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward into the locked position.
5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

6. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode` *`impaired_node_name`*.

7. If automatic giveback was disabled, reenable it:

   `storage failover modify -node local -auto-giveback true`.

8. If AutoSupport is enabled, restore/unsuppress automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVRAM - AFF A70 and AFF A90

Replace the NVRAM in your AFF A70 or AFF A90 system when the non-volatile memory becomes faulty or requires an upgrade. The replacement process involves shutting down the impaired controller, replacing the NVRAM module or the NVRAM DIMM, reassigning the disks, and returning the failed part to NetApp.

The NVRAM module consists of the NVRAM12 hardware and field-replaceable DIMMs.
You can replace a failed NVRAM module or the DIMMs inside the NVRAM module.

Before you begin

- Make sure you have the replacement part available. You must replace the failed component with a replacement component you received from NetApp.
- Make sure all other components in the storage system are functioning properly; if not, contact NetApp Support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the NVRAM module or NVRAM DIMM

Replace the NVRAM module or NVRAM DIMMs using the appropriate following option.

Option 1: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 4/5 in the chassis and follow the specific sequence of steps.

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

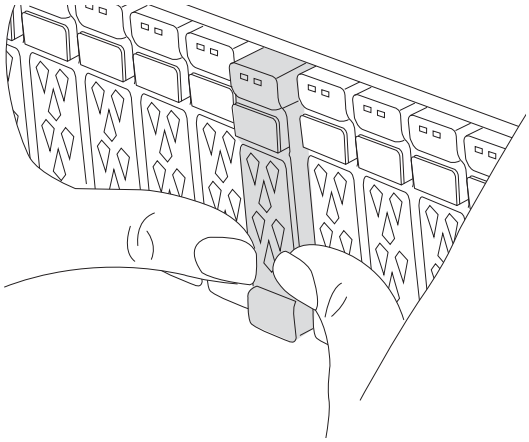

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. If you are not already grounded, properly ground yourself.
4. Unplug the controller's PSUs.

   If your system has DC power, disconnect the power block from the PSUs.

5. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
6. Remove the target NVRAM module from the chassis:
   a. Depress the cam latch button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.
   c. Remove the impaired NVRAM module from the chassis by hooking your finger into the cam lever opening and pulling the module out of the chassis.

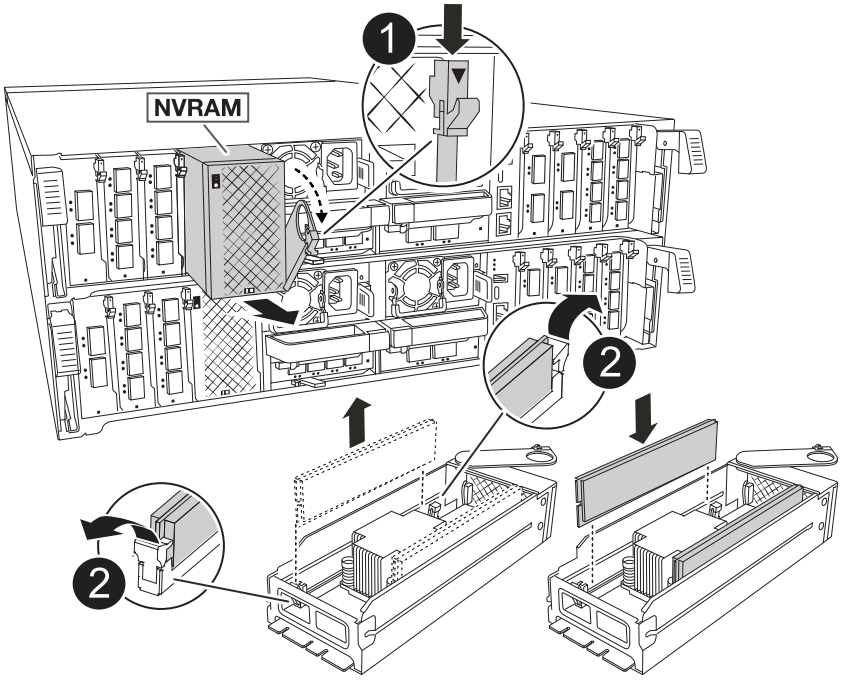

1	Cam locking button
2	DIMM locking tabs

7. Set the NVRAM module on a stable surface.
8. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
9. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 4/5.
   b. Gently slide the module into the slot all the way, and then push the cam latch all the way up to lock the module in place.
10. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

    If you have DC power supplies, reconnect the power block to the power supplies.

11. Rotate the cable management tray up to the closed position.
12. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*.
13. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`.
14. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Option 2: Replace the NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, first must remove the NVRAM module and then replace the target DIMM.

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

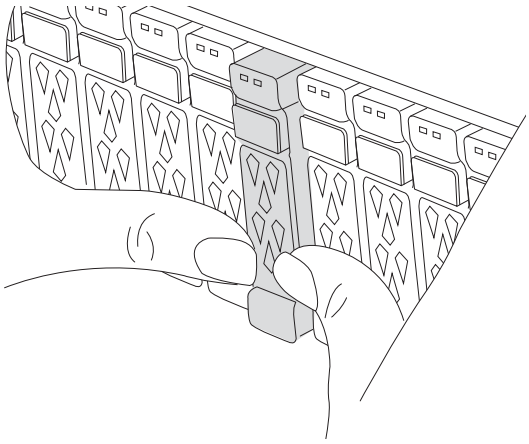

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller's PSUs.

   If your system has DC power, disconnect the power block from the PSUs.

4. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
5. Remove the target NVRAM module from the chassis:
   a. Depress the cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.
   c. Remove the NVRAM module from the chassis by hooking your finger into the cam lever opening and pulling the module out of the chassis.

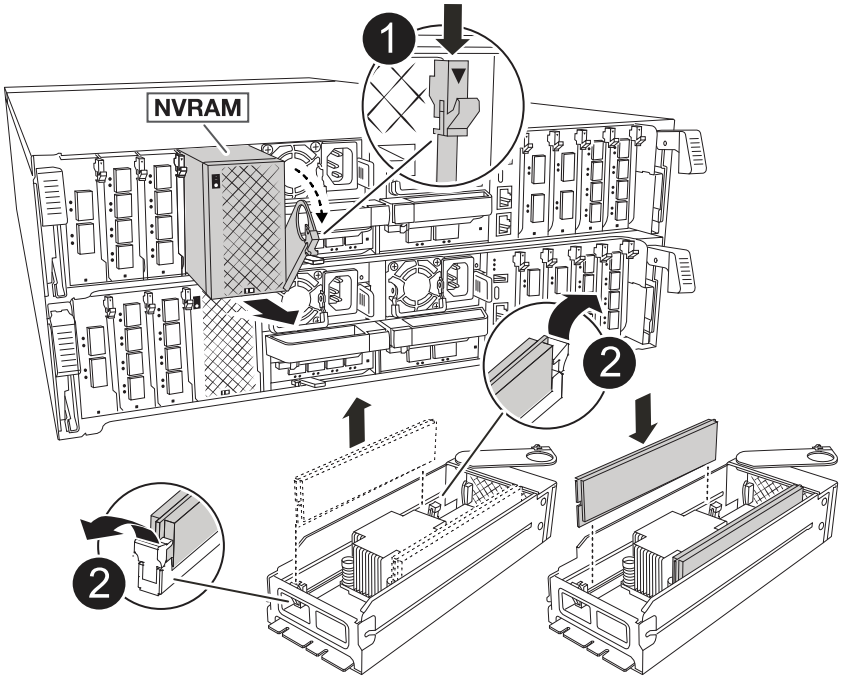

1	Cam locking button
2	DIMM locking tabs

6. Set the NVRAM module on a stable surface.
7. Locate the DIMM to be replaced inside the NVRAM module.

> Consult the FRU map label on the side of the NVRAM module to determine the locations of DIMM slots 1 and 2.

8. Remove the DIMM by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
9. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
10. Install the NVRAM module into the chassis:
    a. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.
11. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

> If you have DC power supplies, reconnect the power block to the power supplies.

12. Rotate the cable management tray up to the closed position.
13. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*.
14. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`.

15. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 3: Reassign disks

You must confirm the system ID change when you boot the controller and then verify that the change was implemented.

Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt:

   `halt`

2. From the LOADER prompt on the controller, boot the controller and enter *y* when prompted to override the system ID due to a system ID mismatch.
3. Wait until the Waiting for giveback message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify that the new partner system ID has been automatically assigned:

   `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node 2 has undergone replacement and has a new system ID of 151759706.

```
node1:> storage failover show
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage:

      `storage failover giveback -ofnode replacement_node_name`

      The controller takes back its storage and completes booting.

If you are prompted to override the system ID due to a system ID mismatch, you should enter *y*.

If the giveback is vetoed, you can consider overriding the vetoes.

For more information, see manual giveback commands to override the veto.

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: *storage failover show*

The output from the `storage failover show` command should not include the System ID changed on partner message.

5. Verify that the disks were assigned correctly:

`storage disk show -ownership`

The disks belonging to the controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 151759706:

```
node1:> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        151759706  151759706  -
151759706 Pool0
1.0.1  aggr0_1  node1 node1           151759706  151759706  -
151759706 Pool0
.
.
.
```

6. If the system is in a MetroCluster configuration, monitor the status of the controller: *metrocluster node show*

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

7. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The controller is the current owner of the disks on the disaster site.

See Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration for more information.

8. If your system is in a MetroCluster configuration, verify that each controller is configured: *metrocluster node show - fields configuration-state*

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

9. Verify that the expected volumes are present for each controller:

   `vol show -node node-name`

10. If storage encryption is enabled, you must restore functionality.
11. Return the impaired controller to normal operation by giving back its storage:

    `storage failover giveback -ofnode impaired_node_name.`

12. If automatic giveback was disabled, reenable it:

    `storage failover modify -node local -auto-giveback true.`

13. If AutoSupport is enabled, restore/unsuppress automatic case creation:

    `system node autosupport invoke -node * -type all -message MAINT=END.`

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NV battery - AFF A70 and AFF A90

Replace the NV battery in your AFF A70 or AFF A90 system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

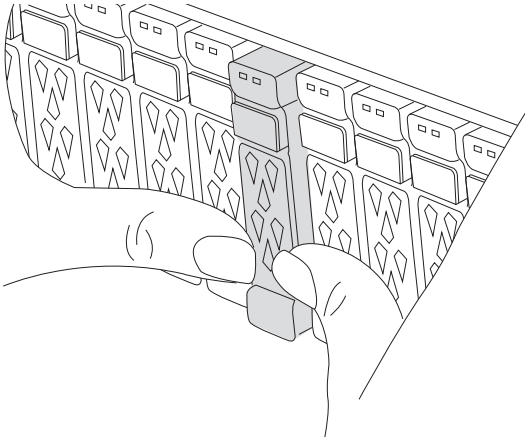

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

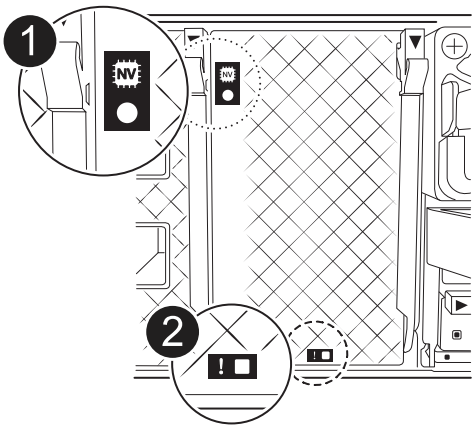

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

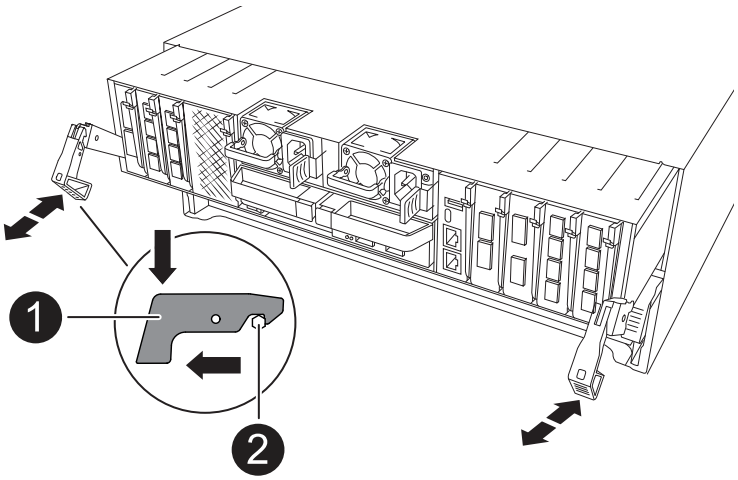

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the NV battery

Remove the failed NV battery from the controller module and install the replacement NV battery.

Steps

1. Open the air duct cover and locate the NV battery.

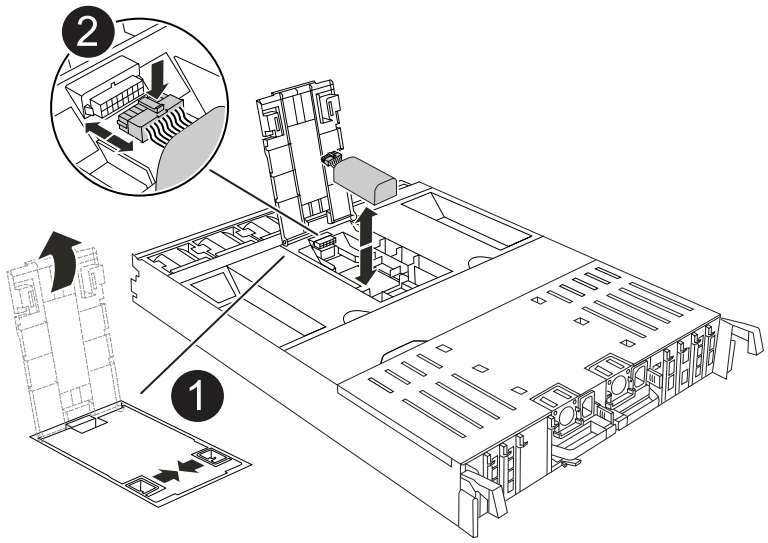

1	NV battery air duct cover
2	NV battery plug

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module, and then set it aside.
5. Remove the replacement battery from its package.
6. Install the replacement battery pack into the controller:
   a. Plug the battery plug into the riser socket and make sure that the plug locks into place.
   b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
7. Close the NV air duct cover.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

   Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward into the locked position.
5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

6. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode` *`impaired_node_name`*.

7. If automatic giveback was disabled, reenable it:

   `storage failover modify -node local -auto-giveback true`.

8. If AutoSupport is enabled, restore/unsuppress automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Overview of add and replace an I/O module - AFF A70 and AFF A90

The AFF A70 and AFF A90 systems offer flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding or replacing an I/O module is

essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your AFF A70 or AFF A90 storage system with the same type of I/O module, or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

- Add an I/O module

  Adding additional modules can improve redundancy, helping to ensure that the system remains operational even if one module fails.

- Replace an I/O module

  Replacing a failing I/O module can restore the system to its optimal operating state.

I/O slot numbering

The I/O slots on AFF A70 and AFF A90 controllers are numbered 1 through 11, as shown in the following illustration.

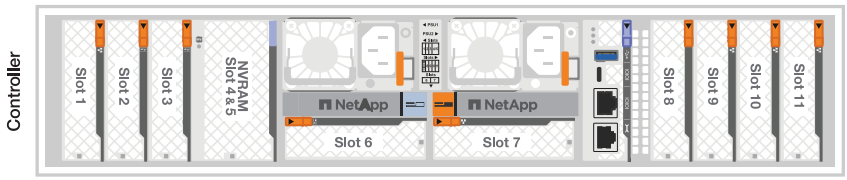

Add an I/O module - AFF A70 and AFF A90

Add an I/O module to your AFF A70 or AFF A90 system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your system when there are empty slots available or when all slots are fully populated.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

A storage system has two location LEDs, one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: MetroCluster configuration

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- Make sure that all other components are functioning properly.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the target slot blanking module from the carrier:
   a. Depress the cam latch on the blanking module in the target slot.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
4. Install the I/O module:
   a. Align the I/O module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
5. Cable the I/O module to the designated device.

   Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

6. Rotate the cable management tray up to the closed position.
7. From the LOADER prompt, reboot the node:

   ```
   bye
   ```

   This reinitializes the I/O module and other components and reboots the node.

8. Give back the controller from the partner controller:

   ```
   storage failover giveback -ofnode target_node_name
   ```

9. Repeat these steps for controller B.
10. From the healthy node, restore automatic giveback if you disabled it:

    ```
    storage failover modify -node local -auto-giveback true
    ```

11. If AutoSupport is enabled, restore automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the target I/O module from the chassis:
   a. Depress the cam latch button.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

      Make sure that you keep track of which slot the I/O module was in.
5. Install the I/O module into the target slot in the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
6. Cable the I/O module to the designated device.
7. Repeat the remove and install steps to replace additional modules for the controller.
8. Rotate the cable management tray up to the closed position.
9. Reboot the controller from the LOADER prompt:_bye_

   This reinitializes the PCIe cards and other components and reboots the node.

   If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement
10. Give back the controller from the partner controller:

    ```
    storage failover giveback -ofnode target_node_name
    ```
11. Enable automatic giveback if it was disabled:

```
storage failover modify -node local -auto-giveback true
```

12. Do one of the following:
    - If you removed a NIC I/O module and installed a new NIC I/O module, use the following network command for each port:

      `storage port modify -node *<node name> -port *<port name> -mode network`

    - If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
13. Repeat these steps for controller B.

Replace an I/O module - AFF A70 and AFF A90

Replace an I/O module in your AFF A70 or AFF A90 system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

You can use this procedure with all versions of ONTAP supported by your storage system.

Before you begin

- You must have the replacement part available.
- Make sure all other components in the storage system are functioning properly; if not, contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace a failed I/O module

To replace an I/O module, locate it within the controller module and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.

   Make sure to label the cables so that you know where they came from.

3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the I/O module from the controller module:

   This following illustration shows removing a horizontal and vertical I/O module. Typically, you will only remove one I/O module.

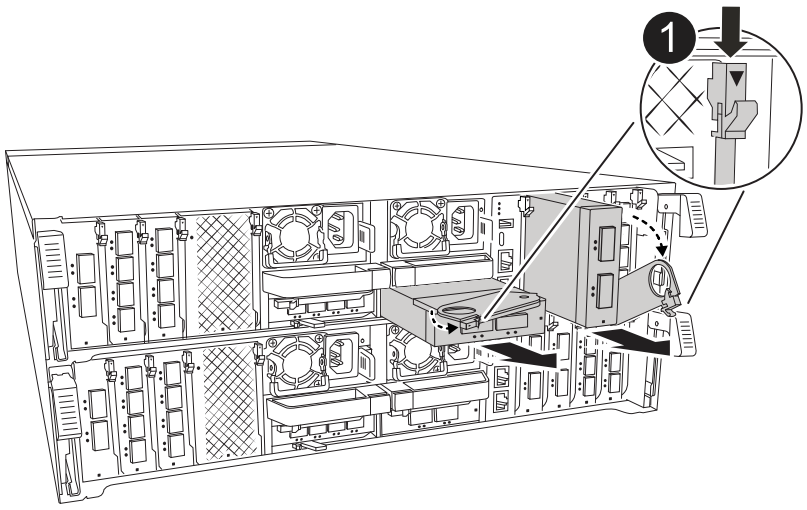

1	Cam locking button

   a. Depress the cam latch button.
   b. Rotate the cam latch do away from the module as far as it will go.
   c. Remove the module from the controller module by hooking your finger into the cam lever opening and

pulling the module out of the controller module.

Make sure that you keep track of which slot the I/O module was in.

5. Set the I/O module aside.
6. Install the replacement I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently slide the module into the slot all the way into the controller module, and then rotate the cam latch all the way up to lock the module in place.
7. Cable the I/O module.
8. Repeat the remove and install steps to replace additional modules for the controller.
9. Rotate the cable management tray into the locked position.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```

3. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```

4. If AutoSupport is enabled, restore automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - AFF A70 and AFF A90

Replace an AC or DC power supply unit (PSU) in your AFF A70 or AFF A90 system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cable, replacing the faulty PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

About this task

This procedure is written for replacing one PSU at a time.

> Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

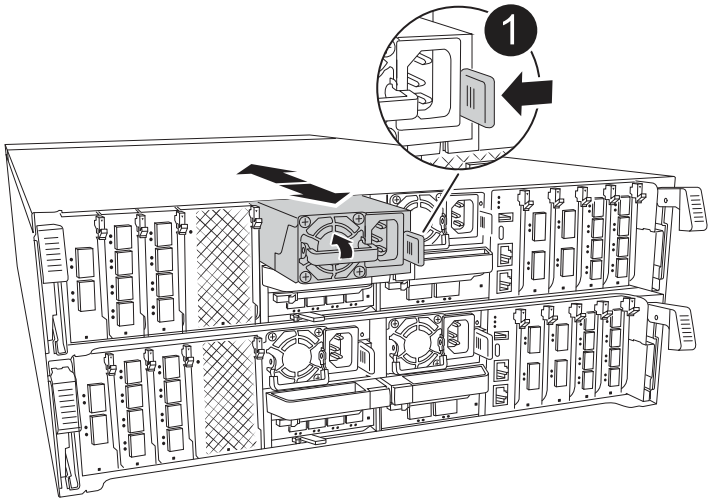

1	Terracotta PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.
6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU.

b. Secure the power cable to the PSU using the power cable retainer.

Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

> The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

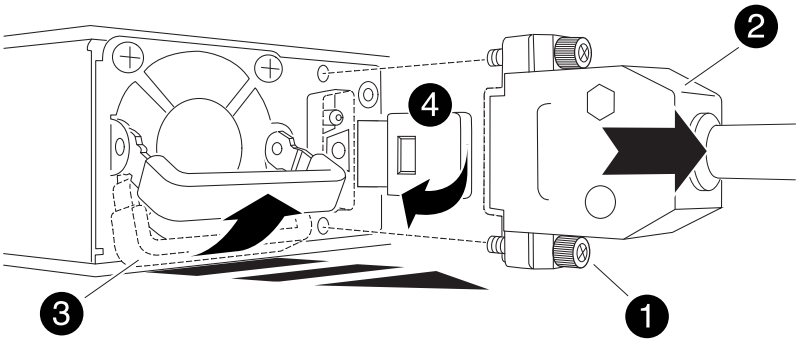

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle
4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.

b. Gently push the PSU into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A70 and AFF A90

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your AFF A70 or AFF A90 system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

- Understand that you can use this procedure with all versions of ONTAP supported by your system.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

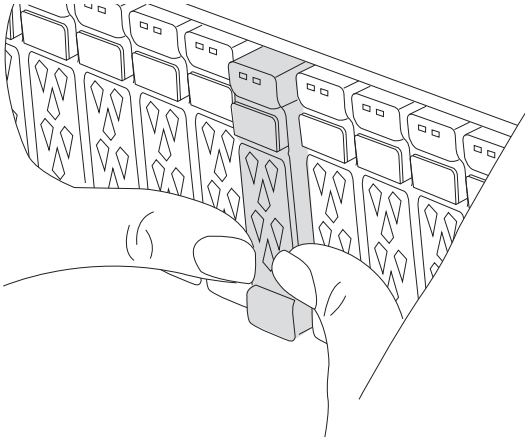

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

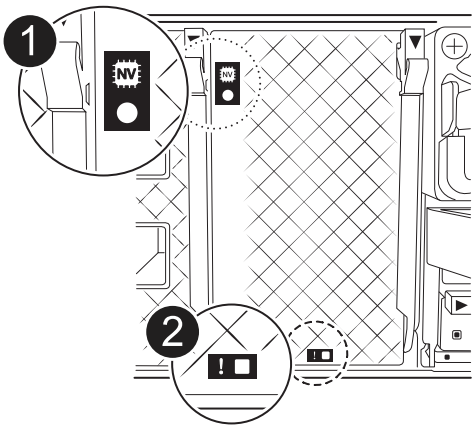

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

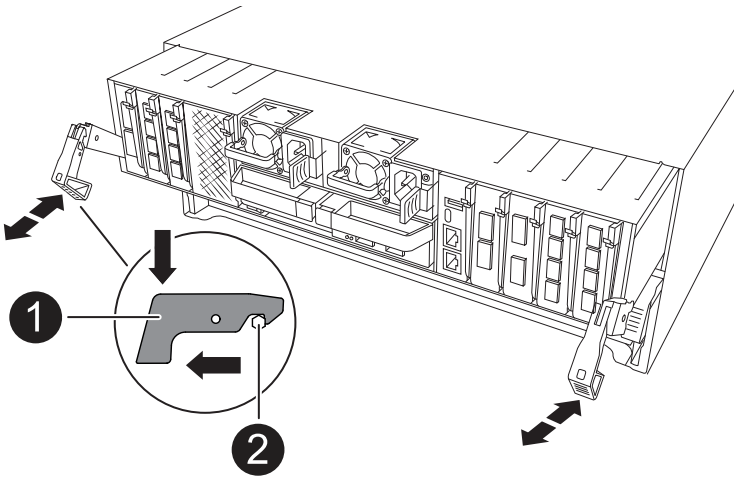

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the RTC battery

Remove the failed RTC battery and install the replacement RTC battery.

Steps

1. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the RTC battery under the air duct.

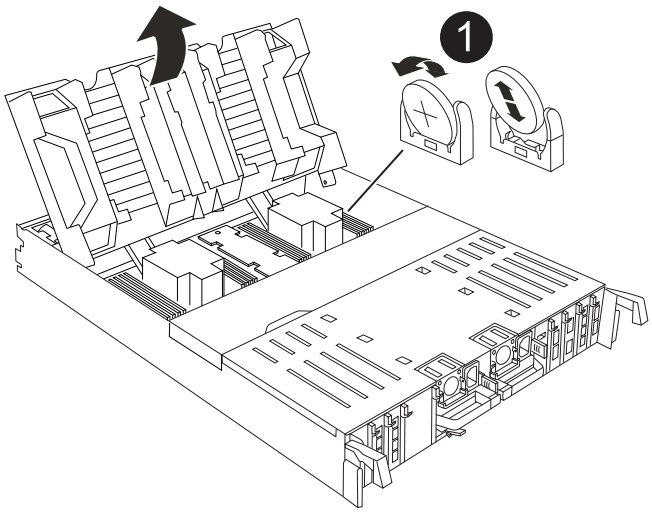

1	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

   Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward into the locked position.
5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

6. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode` *`impaired_node_name`*.

7. If automatic giveback was disabled, reenable it:

   `storage failover modify -node local -auto-giveback true`.

8. If AutoSupport is enabled, restore/unsuppress automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Reset the time and date on the controller

After you replace the RTC battery, insert the controller, and power on for the first BIOS reset, you will see the following error messages:

```
RTC date/time error. Reset date/time to default
```

```
RTC power failure error
```

These messages are expected and you can continue with this procedure.

Steps

1. Check the date and time on the healthy controller with the `cluster date show` command.

   If your system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*

   a. At the LOADER prompt on the target controller, check the time and date with the `cluster date show` command.

   b. If necessary, modify the date with the `set date mm/dd/yyyy` command.

   c. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.

2. Confirm the date and time on the target controller.
3. At the LOADER prompt, enter *bye* to reinitialize the PCIe cards and other components and let the controller reboot.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the system management module - AFF A70 and AFF A90

Replace the System Management module in your AFF A70 or AFF A90 system when it becomes defective or its firmware is corrupted. The replacement process involves shutting down the controller, replacing the failed System Management module, rebooting the controller, updating the license keys, and returning the failed part to NetApp.

The System Management module, located at the back of the controller in slot 8, contains onboard components for system management, as well as ports for external management. The target controller must be shut down to replace an impaired System Management module or to replace the boot media.

Before you begin

- Make sure all other system components are working properly.
- Make sure that the partner controller is able to take over the impaired controller.
- Make sure you replace the failed component with a replacement component you received from NetApp.

About this task

This procedure uses the following terminology:

- The impaired controller is the controller on which you are performing maintenance.
- The healthy controller is the HA partner of the impaired controller.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the System Management module

Replace the impaired system management module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

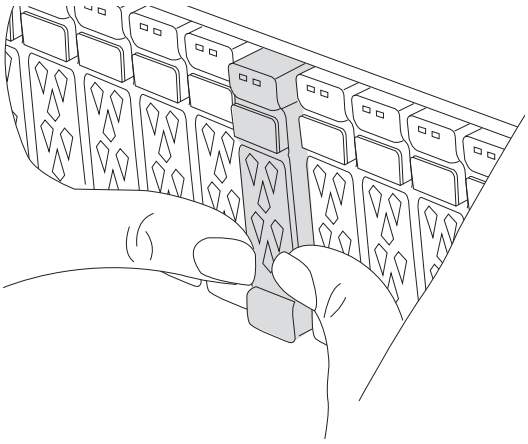

2. Make sure NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged. If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

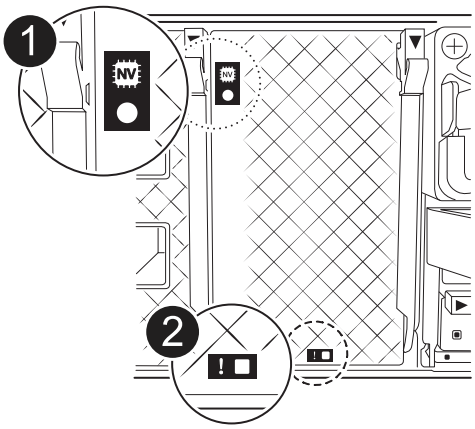

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
4. Unplug the controller's PSUs.

> If your system has DC power, disconnect the power block from the PSUs.

5. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
6. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.

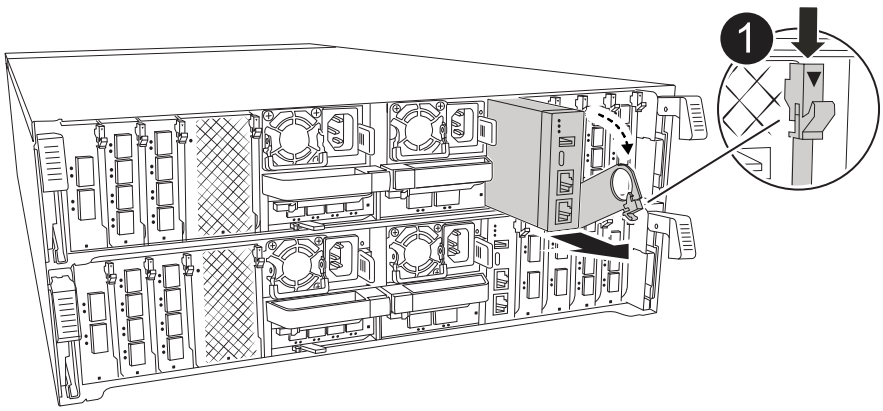

1	System Management module cam latch

7. Remove the System Management module:
   a. Depress the system management cam button.
   The cam lever moves away from the chassis.
   b. Rotate the cam lever all the way down.
   c. Loop your finger into the cam lever and pull the module straight out of the system.
   d. Place the System Management module on an anti-static mat, so that the boot media is accessible.
8. Move the boot media to the replacement System Management module:

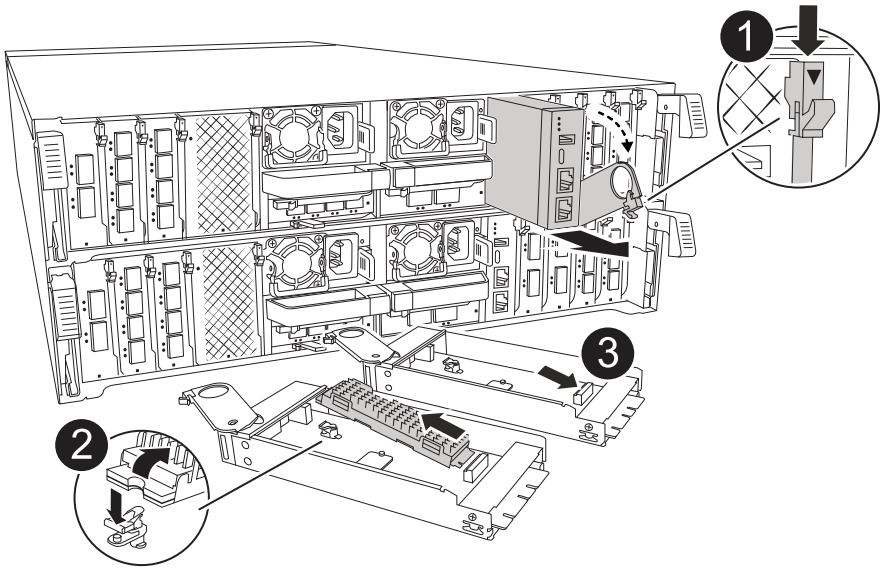

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue locking button.
The boot media rotates slightly upward.

b. Rotate the boot media up, slide it out of the socket.

c. Install the boot media in the replacement System Management module:

i. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

ii. Rotate the boot media down toward until it engages the locking button. Depress the blue locking if necessary.

9. Install the system management module:

a. Align the edges of the replacement System Management module with the system opening and gently push it into the controller module.

b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

10. Recable the System Management module.

11. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

If you have DC power supplies, reconnect the power block to the power supplies.

12. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

Reboot the controller module.

Steps

1. Enter *bye* at the LOADER prompt.
2. Return the controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

3. Restore automatic giveback:

   `storage failover modify -node local -auto-giveback true`

4. If an AutoSupport maintenance window was triggered, end it:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 4: Install licenses and register serial number

You must install new licenses for the node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the node. However, if the node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the for the node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`

3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`
4. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF A20, AFF A30, and AFF A50 systems

Install and setup

Installation and setup workflow - AFF A20, AFF A30, and AFF A50

To install and set up your AFF A20, AFF A30, or AFF A50 storage system, you must review the installation requirements, prepare your site, install and cable the hardware components, power on the storage system, and set up the ONTAP cluster.

Review the installation requirements

Before installing your storage system, it must meet the installation requirements.

Prepare for installation

To prepare for installation, get the site ready, check environmental and electrical requirements, and ensure there's enough rack space. Then, unpack the equipment, compare contents to the packing slip, and register the hardware to access support benefits.

Install the hardware

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system and shelves in the cabinet or telco rack.

Cable the hardware

To cable the hardware, connect the controllers to your network and then to your shelves.

Power on the storage system

To power on your storage system, power on each shelf and assign a unique shelf ID as needed, and then power on the controllers.

6 Set up your cluster

After you've powered on your storage system, you set up your cluster.

Installation requirements - AFF A20, AFF A30, and AFF A50

Review the requirements for your AFF A20, AFF A30, or AFF A50 storage system.

Equipment needed for install

To install your storage system, you need the following equipment and tools.

- Access to a Web browser to configure your storage system
- Electrostatic discharge (ESD) strap
- Flashlight
- Laptop or console with a USB/serial connection
- Phillips #2 screwdriver

Lifting precautions

Storage systems and shelves are heavy. Exercise caution when lifting and moving these items.

Storage system weight

Take the necessary precautions when moving or lifting your storage system.

An A1K storage system can weigh up to 62.83 lbs (28.5 kg). To lift the storage system, use two people or a hydraulic lift.

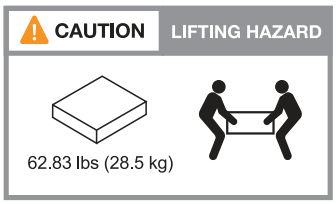

The storage system can weigh up to 61.5 lbs (27.9 kg). To lift the storage system, use two people or a hydraulic lift.

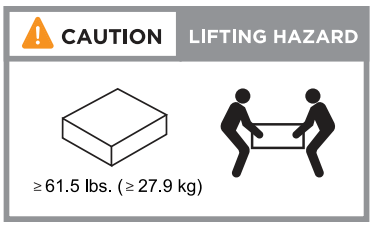

Shelf weight

Take the necessary precautions when moving or lifting your shelf.

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

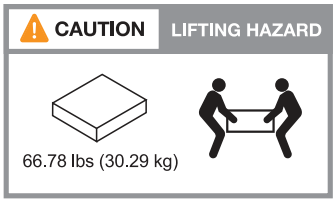

An NS224 shelf with NSM100B modules can weigh up to 56.8 lbs (25.8 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

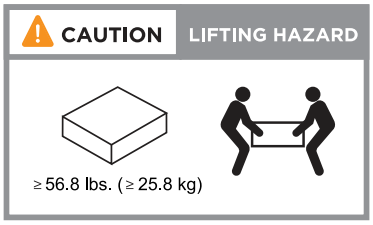

Related information

- Safety information and regulatory notices

What's next?

After you've reviewed the installation requirements and considerations for your storage system, you prepare for installation.

Prepare to install - AFF A20, AFF A30, and AFF A50

Prepare to install your AFF A20, AFF A30, or AFF A50 storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the storage system to access support benefits.

Step 1: Prepare the site

To install your storage system, ensure that the site and the cabinet or rack that you plan to use meet specifications for your configuration.

Steps

1. Use NetApp Hardware Universe to confirm that your site meets the environmental and electrical requirements for your storage system.
2. Make sure you have adequate cabinet or rack space for your storage system, shelves, and any switches:
   - 4U in an HA configuration
   - 2U for each NS224 storage shelf
   - 2U for a storage system
   - 2U for each NS224 storage shelf
   - 1U for most switches
3. Install any required network switches.

   See the Switch documentation for installation instructions and NetApp Hardware Universe for compatibility information.

Step 2: Unpack the boxes

After you've ensured that the site and the cabinet or rack that you plan to use for your storage system meet the required specifications, unpack all boxes and compare the contents to the items on the packing slip.

Steps

1. Carefully open all the boxes and lay out the contents in an organized manner.
2. Compare the contents you've unpacked with the list on the packing slip.

   You can get your packing list by scanning the QR code on the side of the shipping carton.

   The following items are some of the contents you might see in the boxes.

   Ensure that everything in the boxes matches the list on the packing slip. If there are any discrepancies, note them down for further action.

Hardware

- Bezel
- Cable management device
- Storage system
- Rail kits with instructions (optional)
- Storage shelf (if you ordered additional storage)

Cables

- Management Ethernet cables (RJ-45 cables)
- Network cables
- Power cords
- Storage cables (if you ordered additional storage)
- USB-C serial console cable

Step 3: Register your storage system

After you've ensured that your site meets the requirements for your storage system specifications, and you've verified that you have all the parts you ordered, you should register your storage system.

Steps

1. Locate the System Serial Numbers (SSN) for every controller being installed. You can find the serial numbers in the following locations:
2. You can find the serial numbers in the following locations:
   - On the packing slip
   - In your confirmation email
   - On each controller's System Management module
   - On each controller

3. Go to the NetApp Support Site.
4. Determine whether you need to register your storage system:

If you are a…	Follow these steps…
Existing NetApp customer	a. Sign in with your username and password. b. Select **Systems** > **My Systems**. c. Confirm that the new serial numbers are listed. d. If it is not, follow the instructions for new NetApp customers.
New NetApp customer	a. Click **Register Now**, and create an account. b. Select **Systems** > **Register Systems**. c. Enter the storage system's serial numbers and requested details. After your registration is approved, you can download any required software. The approval process might take up to 24 hours.

What's next?

After you've prepared to install your storage system, you install the hardware for your storage system.

Install the hardware - AFF A20, AFF A30, and AFF A50

After you prepare to install your AFF A20, AFF A30, or AFF A50 storage system, install the hardware for the storage system. First, install the rail kits. Then install and secure your storage system in a cabinet or telco rack.

Skip this step if your storage system came in a cabinet.

Before you begin

- Make sure you have the instructions packaged with the rail kit.
- Be aware of the safety concerns associated with the weight of the storage system and shelf.
- Understand that the airflow through the storage system enters from the front where the bezel or end caps are installed and exhausts out the rear where the ports are located.

Steps

1. Install the rail kits for your storage system and shelves as needed, using the instructions included with the kits.
2. Install and secure your storage system in the cabinet or telco rack:
   a. Position the storage system onto the rails in the middle of the cabinet or telco rack, and then support the storage system from the bottom and slide it into place.
   b. Make sure that the guiding pins of the cabinet or telco rack are securley in the chassis guide slots.
   c. Secure the storage system to the cabinet or telco rack using the included mounting screws.
3. Attach the bezel to the front of the storage system.
4. Attach the cable management devices to the rear of the storage system.

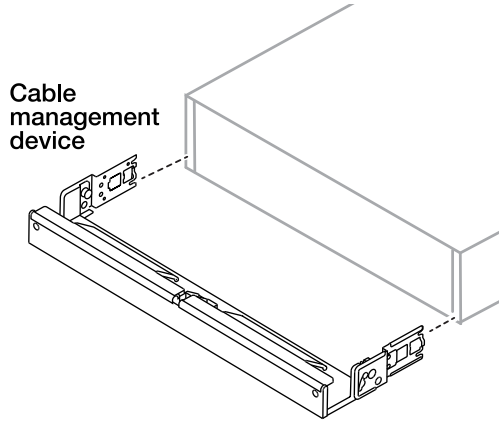

5. Install and secure the shelf as needed.
   a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

      If you are installing multiple shelves, place the first shelf directly above the controllers. Place the second shelf directly under the controllers. Repeat this pattern for any additional shelves.

   b. Secure the shelf to the cabinet or telco rack using the included mounting screws.

      What's next?

      After you've installed the hardware for your storage system, you cable the hardware.

Cable the hardware - AFF A20, AFF A30, and AFF A50

After you install your AFF A20, AFF A30, or AFF A50 storage system hardware, cable the controllers to the network and shelves.

Before you begin

Contact your network administrator for information about connecting the storage system to your network switches.

About this task

- The cabling graphics have arrow icons showing the proper orientation (up or down) of the cable connector pull-tab when inserting a connector into a port.

  As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it over and try again.

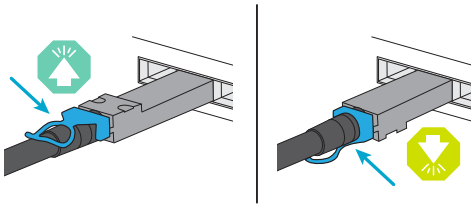

- If cabling to an optical switch, insert the optical transceiver into the controller port before cabling to the switch port.

Step 1: Cable the cluster/HA connections

Create the ONTAP cluster connections. For switchless clusters, connect the controllers to each other. For switched clusters, connect the controllers to the cluster network switches.

The cluster/HA cabling examples show common configurations.

If you do not see your configuration here, go to NetApp Hardware Universe for comprehensive configuration and slot priority information to cable your storage system.

Switchless cluster cabling

AFF A30 or AFF A50 with two 2-port 40/100 GbE I/O modules

Steps

1. Cable the Cluster/HA interconnect connections:

> The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O modules in slots 2 and 4). The ports are 40/100 GbE.

a. Cable controller A port e2a to controller B port e2a.

b. Cable controller A port e4a to controller B port e4a.

> I/O module ports e2b and e4b are unused and available for host network connectivity.

100 GbE Cluster/HA interconnect cables

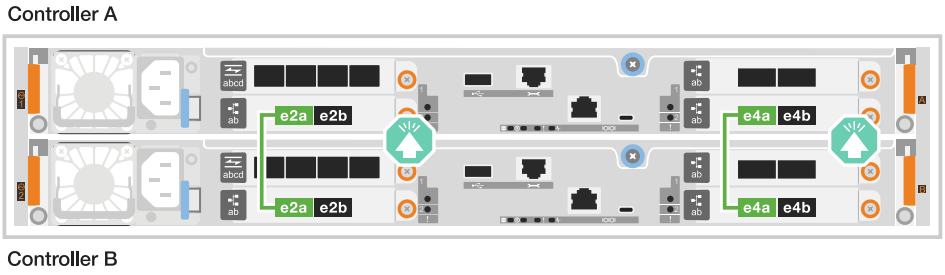

AFF A30 or AFF A50 with one 2-port 40/100 GbE I/O module

Steps

1. Cable the Cluster/HA interconnect connections:

> The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

a. Cable controller A port e4a to controller B port e4a.

b. Cable controller A port e4b to controller B port e4b.

100 GbE Cluster/HA interconnect cables

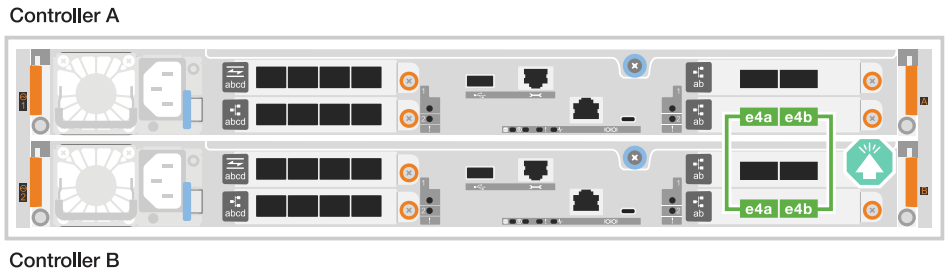

AFF A20 with one 2-port 10/25 GbE I/O module

Steps

1. Cable the Cluster/HA interconnect connections:

> The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 10/25 GbE.

a. Cable controller A port e4a to controller B port e4a.

b. Cable controller A port e4b to controller B port e4b.

25 GbE Cluster/HA interconnect cables

Switched cluster cabling

AFF A30 or AFF A50 with two 2-port 40/100 GbE I/O modules

Steps

1. Cable the Cluster/HA interconnect connections:

> The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O modules in slots 2 and 4). The ports are 40/100 GbE.

a. Cable controller A port e4a to cluster network switch A.

b. Cable controller A port e2a to cluster network switch B.

c. Cable controller B port e4a to cluster network switch A.

d. Cable controller B port e2a to cluster network switch B.

> I/O module ports e2b and e4b are unused and available for host network connectivity.

40/100 GbE Cluster/HA interconnect cables

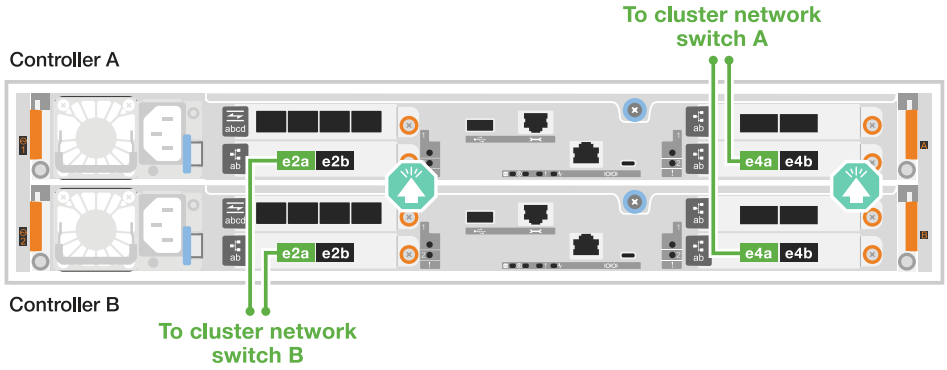

AFF A30 or AFF A50 with one 2-port 40/100 GbE I/O module

Steps

1. Cable the controllers to the cluster network switches:

> The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

a. Cable controller A port e4a to cluster network switch A.

b. Cable controller A port e4b to cluster network switch B.

c. Cable controller B port e4a to cluster network switch A.

d. Cable controller B port e4b to cluster network switch B.

40/100 GbE Cluster/HA interconnect cables

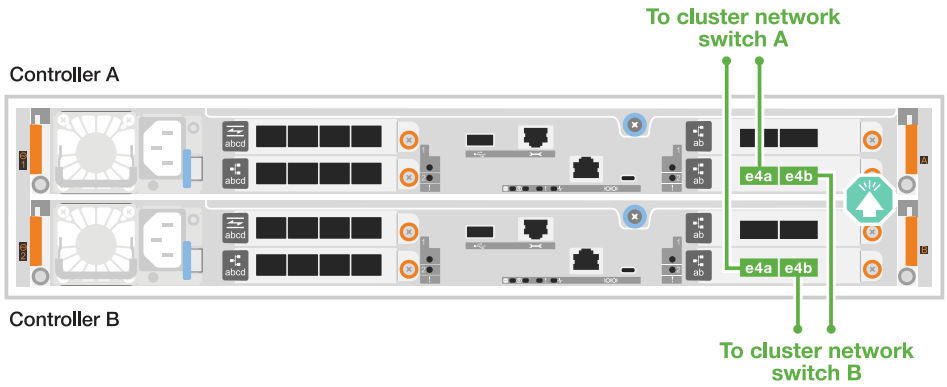

AFF A20 with one 2-port 10/25 GbE I/O module

1. Cable the controllers to the cluster network switches:

   The cluster interconnect traffic and the HA traffic share the same physical ports(on the I/O module in slot 4). The ports are 10/25 GbE.

   a. Cable controller A port e4a to cluster network switch A.
   b. Cable controller A port e4b to cluster network switch B.
   c. Cable controller B port e4a to cluster network switch A.
   d. Cable controller B port e4b to cluster network switch B.

10/25 GbE Cluster/HA interconnect cables

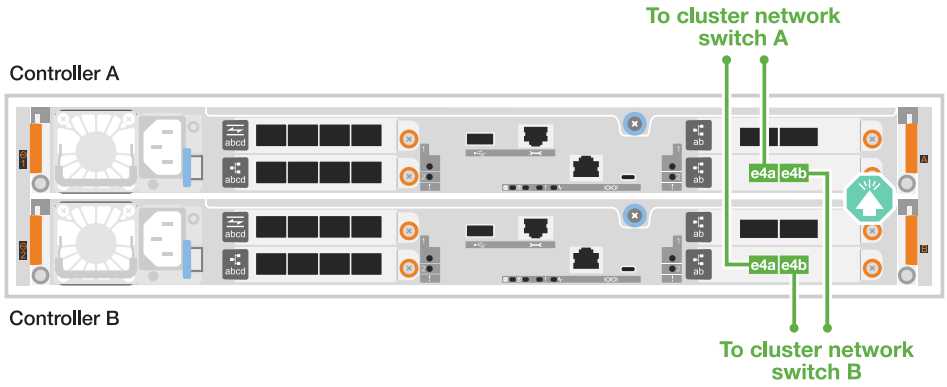

Step 2: Cable the host network connections

Cable the controllers to your Ethernet or FC host network.

The host network cabling examples show common configurations.

If you do not see your configuration here, go to NetApp Hardware Universe for comprehensive configuration and slot priority information to cable your storage system.

Ethernet host cabling

AFF A30 or AFF A50 with two 2-port 40/100 GbE I/O modules

Steps

1. On each controller, cable ports e2b and e4b to the Ethernet host network switches.

> The ports on I/O modules in slot 2 and 4 are 40/100 GbE (host connectivity is 40/100 GbE).

40/100 GbE cables

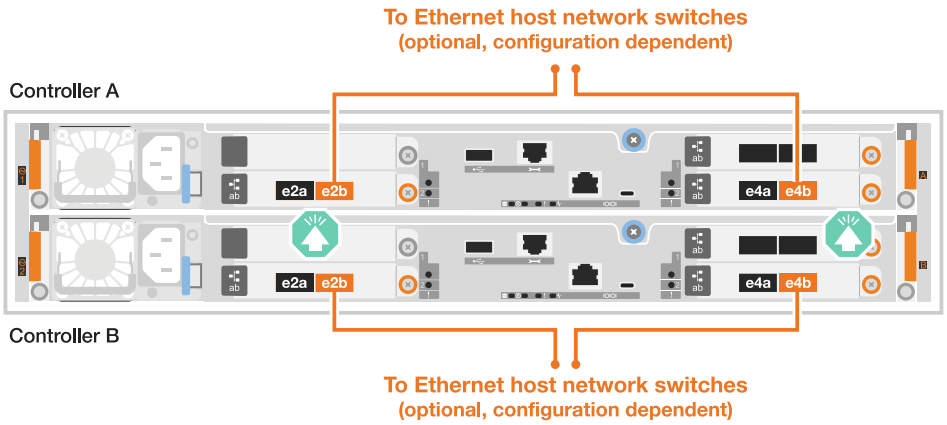

AFF A20, A30 or AFF A50 with one 4-port 10/25 GbE I/O module

Steps

1. On each controller, cable ports e2a, e2b, e2c and e2d to the Ethernet host network switches.

10/25 GbE cables

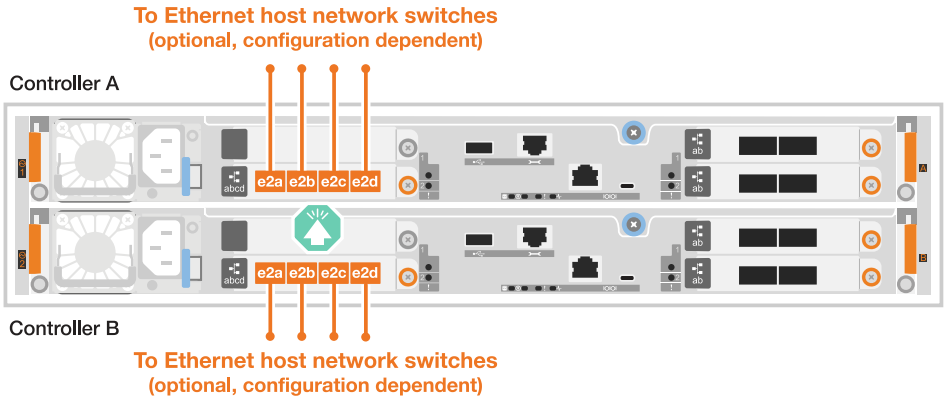

FC host cabling

AFF A20, A30 or AFF A50 with one 4-port 64 Gb/s FC I/O module

Steps

1. On each controller, cable ports 1a, 1b, 1c and 1d to the FC host network switches.

64 Gb/s FC cables

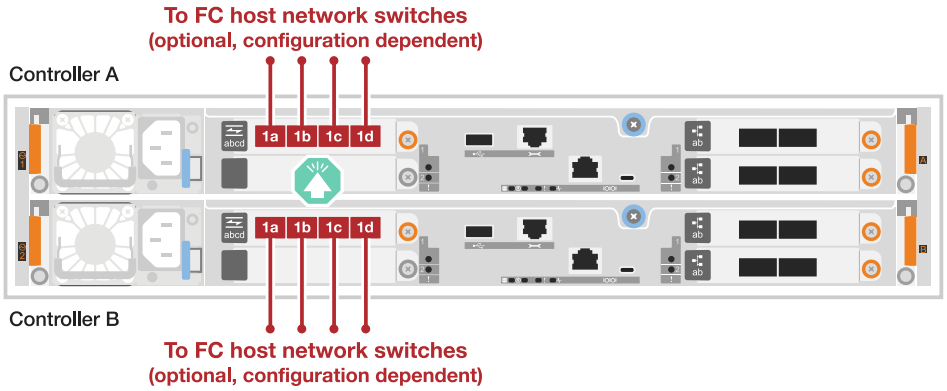

Step 3: Cable the management network connections

Cable the controllers to your management network.

1. Cable the management (wrench) ports on each controller to the management network switches.

   1000BASE-T RJ-45 cables

   

   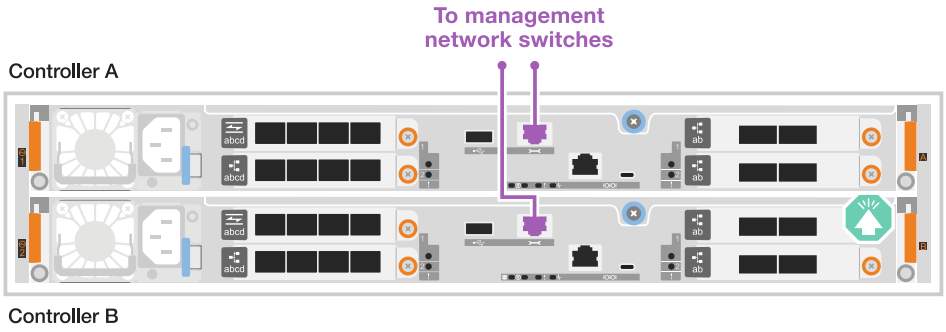

   Do not plug in the power cords yet.

Step 4: Cable the shelf connections

This procedure shows you how to cable the controllers to one NS224 shelf.

About this task

- For the maximum number of shelves supported for your storage system and for all of your cabling options, such as optical and switch-attached, see NetApp Hardware Universe.
- The NS224 shelf cabling procedure shows NSM100B modules instead of NSM100 modules. The cabling is the same regardless of the type of NSM modules used, only the port names are different:
  - NSM100B modules use ports e1a and e1b on an I/O module in slot 1.
  - NSM100 modules use built-in (onboard) ports e0a and e0b.
- You cable each controller to each NSM module on the NS224 shelf using the storage cables that came with your storage system, which could be the following cable type:

  100 GbE QSFP28 copper cables

  

- The graphics show controller A cabling in blue and controller B cabling in yellow.

Steps

1. Cable controller A to the shelf:
   a. Cable controller A port e3a to NSM A port e1a.
   b. Cable controller A port e3b to NSM B port e1b.

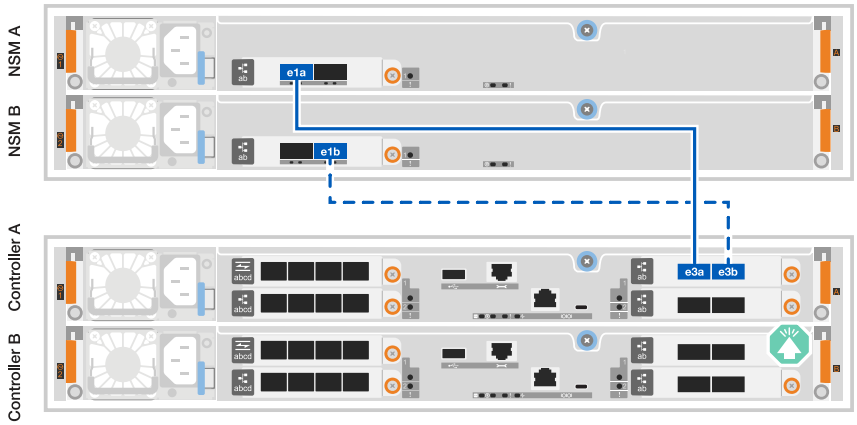

2. Cable controller B to the shelf:
   a. Cable controller B port e3a to NSM B port e1a.
   b. Cable controller B port e3b to NSM A port e1b.

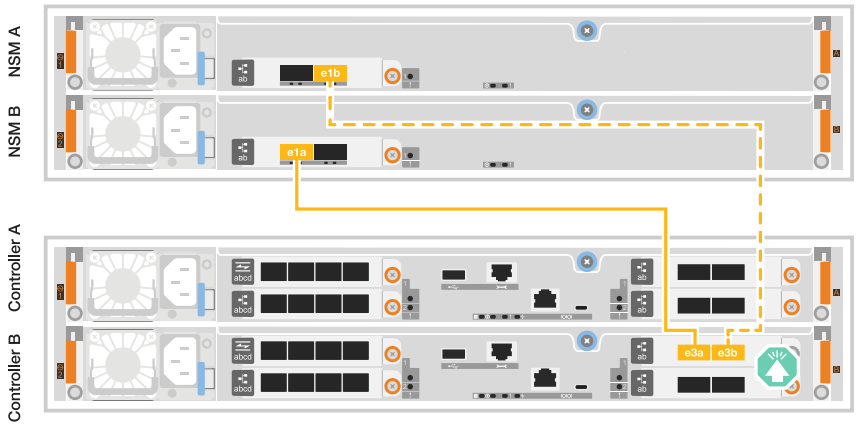

What's next?

After you've cabled the hardware for your storage system, you power on the storage system.

Power on the storage system - AFF A20, AFF A30, and AFF A50

After you cable the controllers to the network and shelves in your AFF A20, AFF A30, or AFF A50 storage system, you power on your shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

Each shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup.

Before you begin

Make sure you have a paperclip or narrow tipped ball point pen for setting NS224 storage shelf IDs.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed

shelf ID of 00.

- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, and then plug them back in) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, and then connecting the power cords to power sources on different circuits.

   The shelf powers on and boots automatically when plugged into the power source.

2. Remove the left end cap to access the shelf ID button behind the faceplate.

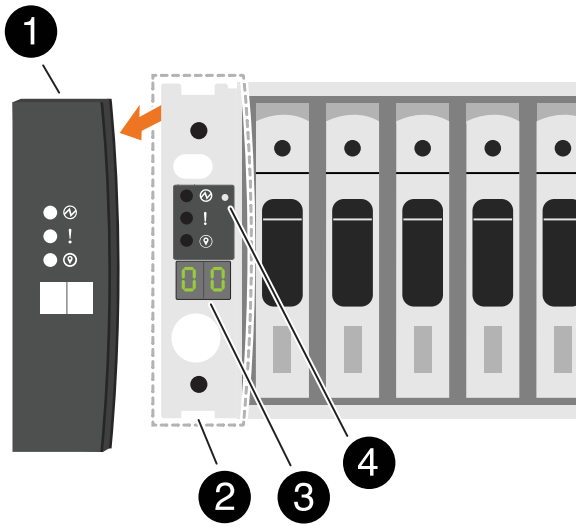

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID number
4	Shelf ID button

3. Change the first number of the shelf ID:
   a. Insert the straightened end of a paperclip or narrow tipped ball point pen into the small hole to press the shelf ID button.
   b. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

      It can take up to 15 seconds for the number to blink. This activates the shelf ID programming mode.

      If the ID takes longer than 15 seconds to blink, press and hold the shelf ID button again, making sure to press it in all the way.

   c. Press and release the shelf ID button to advance the number until you reach the desired number from

0 to 9.

Each press and release duration can be as short as one second.

The first number continues to blink.

4. Change the second number of the shelf ID:
   a. Press and hold the button until the second number on the digital display blinks.

      It can take up to three seconds for the number to blink.

      The first number on the digital display stops blinking.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

      The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the shelf ID button until the second number stops blinking.

   It can take up to three seconds for the number to stop blinking.

   Both numbers on the digital display start blinking and the amber LED illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.
   a. Unplug the power cord from both power supplies on the shelf.
   b. Wait 10 seconds.
   c. Plug the power cords back into the shelf power supplies to complete the power cycle.

      A power supply is powered on as soon as the power cord is plugged in. Its bicolored LED should illuminate green.

7. Replace the left end cap.

Step 2: Power on the controllers

After you've powered on your shelves and assigned them unique IDs, power on the storage controllers.

Steps

1. Connect your laptop to the serial console port. This will allow you to monitor the boot sequence when the controllers are powered on.
   a. Set the serial console port on the laptop to 115,200 baud with N-8-1.

      

      See your laptop's online help for instructions on how to configure the serial console port.

   b. Using the console cable provided with your storage system, connect one end of the console cable to your laptop and the other end to the serial console port on controller A.
   c. Connect the laptop to the switch on the management subnet.

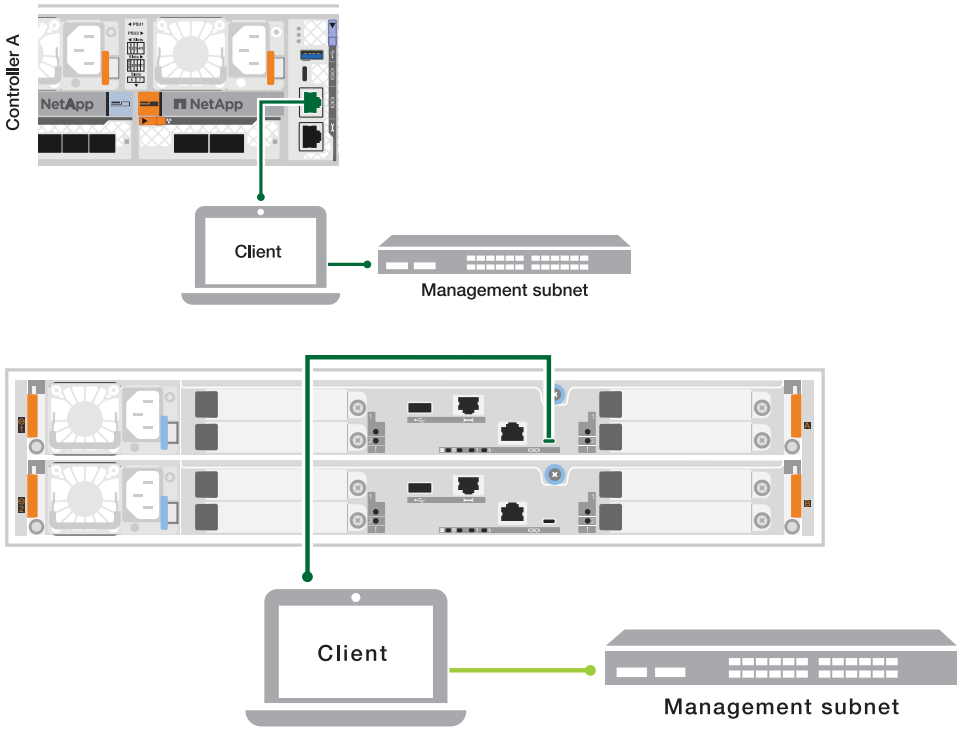

2. Assign a TCP/IP address to the laptop, using one that is on the management subnet.
3. Plug the two power cords into the controller power supplies, and then connect them to power sources on different circuits.

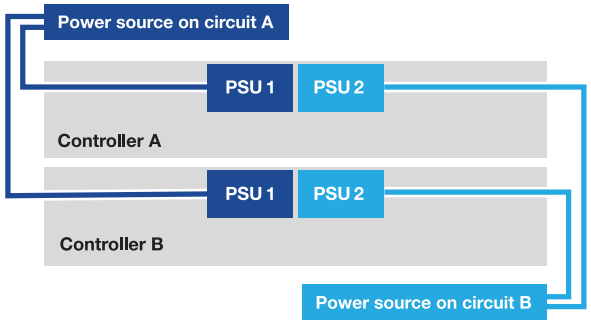

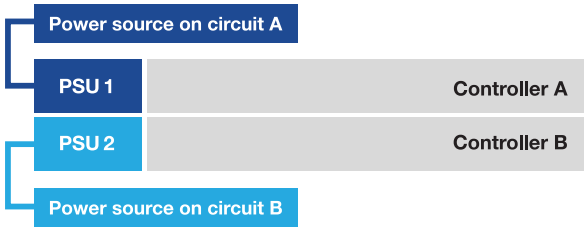

   - The system begins to boot. Initial booting might take up to eight minutes.
   - The LEDs flash on and the fans start, which indicates that the controllers are powering on.
   - The fans might be very noisy when they first start up. The fan noise during start-up is normal.
   - The shelf ID display on the front of the system chassis does not illuminate.
4. Secure the power cords using the securing device on each power supply.

What's next?

After you've powered on your storage system, you set up your cluster.

Maintain

Overview of the maintenance procedures - AFF A20, AFF A30, and AFF A50

Maintain the hardware of your AFF A20, AFF A30, or AFF A50 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the AFF A20, AFF A30, or AFF A50 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF A20, AFF A30, or AFF A50 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media- manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot the image from a USB drive and restore the configuration from the partner node.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller and drives.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.

NV battery	The non-volatile memory (NV) battery is responsible for providing power to the NVMEM components while data in-flight is being destaged to flash memory after a power loss.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - AFF A20, AFF A30, and AFF A50

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your AFF A20, AFF A30, or AFF A50 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

Review the boot media requirements

Review the requirements for boot media replacement.

Shut down the controller

Shut down the controller in your storage system when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the impaired controller and install the replacement boot media.

Restore the image on the boot media

Restore the ONTAP image from the partner controller.

Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - AFF A20, AFF A30, and AFF A50

Before replacing the boot media in your AFF A20, AFF A30, and AFF A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0M (wrench) port on the impaired controller is working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Review the following requirements.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - AFF A20, AFF A30, and AFF A50

Shut down the impaired controller in your AFF A20, AFF A30, or AFF A50 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`)

for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   `cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - AFF A20, AFF A30, and AFF A50

The boot media in your AFF A20, AFF A30, or AFF A50 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot

media, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

About this task

If needed, you can turn on the platform chassis location (blue) LEDs to aid in physically locating the affected platform. Log into the BMC using SSH and enter the `system location-led on` command.

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

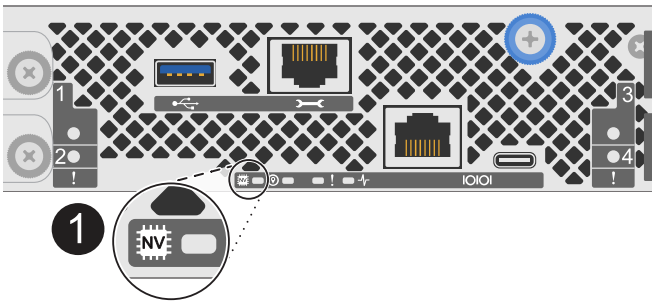

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

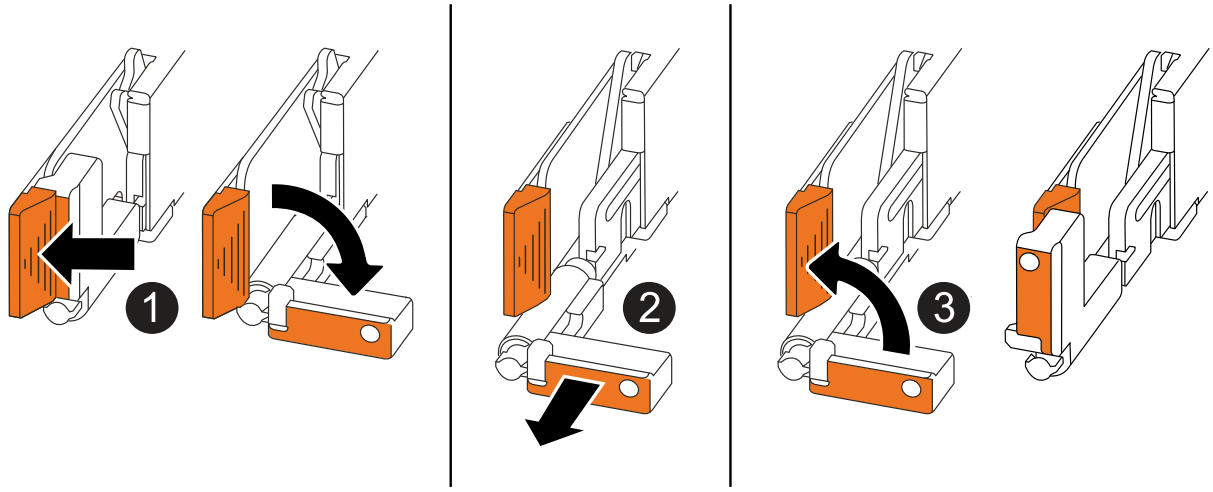

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.

7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:

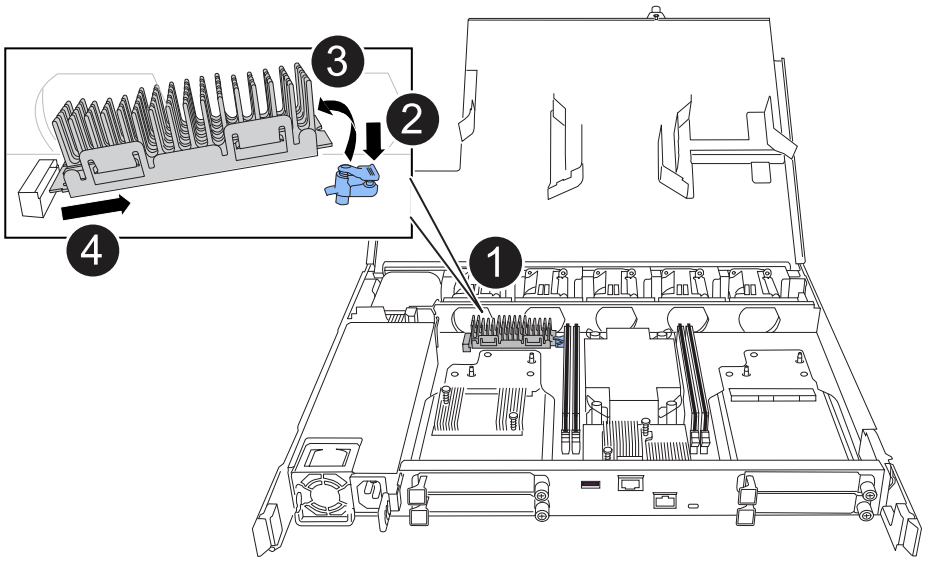

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:
   a. Remove the boot media from its package.
   b. Slide the socket end of the boot media into its socket.
   c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

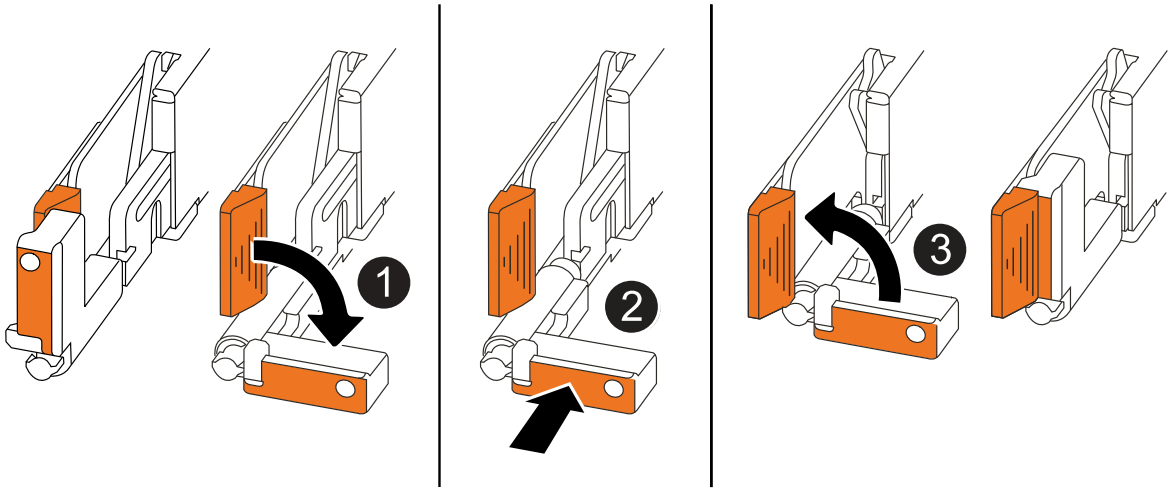

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.

   Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

      The controller boots to the LOADER prompt when fully seated in the chassis. It gets its power from the partner controller.

b. Rotate the controller handles up and lock in place with the tabs.

5. Reconnect the power cord to the PSU on the impaired controller.

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - AFF A20, AFF A30, and AFF A50

After installing the new boot media device in your AFF A20, AFF A30, or AFF A50 storage system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media part to NetApp - AFF A20, AFF A30, and AFF A50

If a component in your AFF A20, AFF A30, or AFF A50 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF A20, AFF A30, and AFF A50

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the AFF A20, AFF A30, or AFF A50 storage system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

If your system is running in ONTAP 9.17.1 and later, use the automatic boot recovery procedure.

To get started, review the recovery requirements, shut down the controller, replace the boot media, use the USB drive to restore the image, and reapply encryption settings if necessary.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check onboard encryption keys

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the impaired controller and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONTAP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - AFF A20, AFF A30, and AFF A50

Before replacing the boot media in your AFF A20, AFF A30, or AFF A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption support for manual boot media recovery - AFF A20, AFF A30, and AFF A50

To ensure data security on your A20, AFF A30, or AFF A50 storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

What's next?

After checking the encryption key support and status on the boot media, you need to shut down the controller.

Shut down the controller for manual boot media recovery - AFF A20, AFF A30, and AFF A50

Shut down the impaired controller in your AFF A20, AFF A30, or AFF A50 storage system to prevent data loss and maintain system stability during the manual boot media recovery process.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After shutting down the controller, you need to replace the boot media.

Replace the boot media and prepare for manual boot recovery - AFF A20, AFF A30, and AFF A50

The boot media in your AFF A20, AFF A30, or AFF A50 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

About this task

If needed, you can turn on the platform chassis location (blue) LEDs to aid in physically locating the affected platform. Log into the BMC using SSH and enter the `system location-led` *`on`* command.

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led` *`off`* command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.

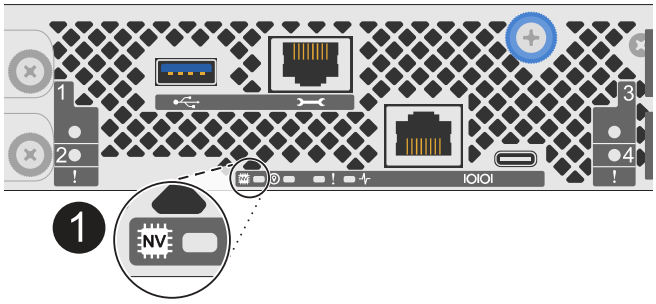

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

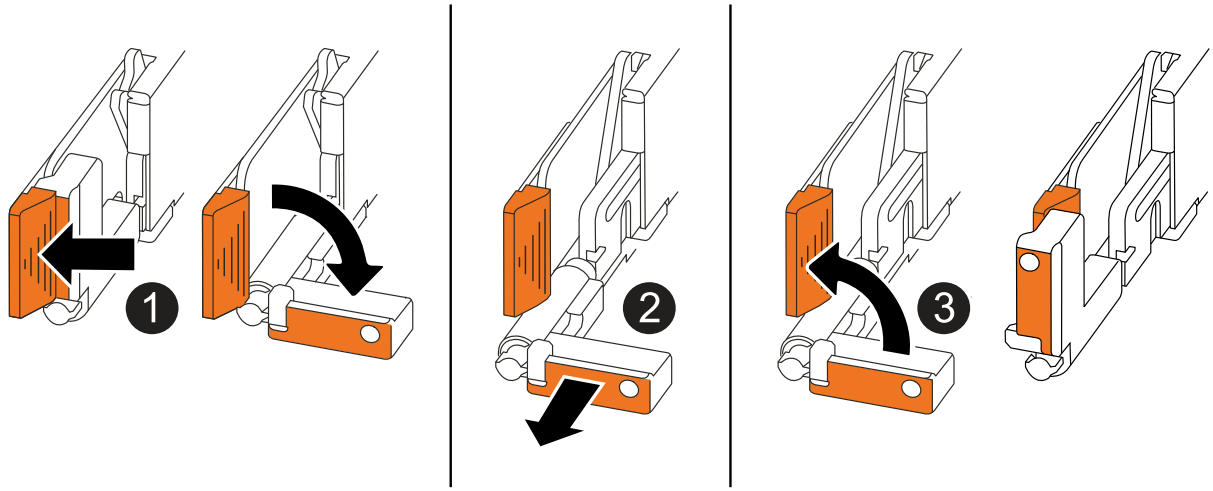

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.
7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:

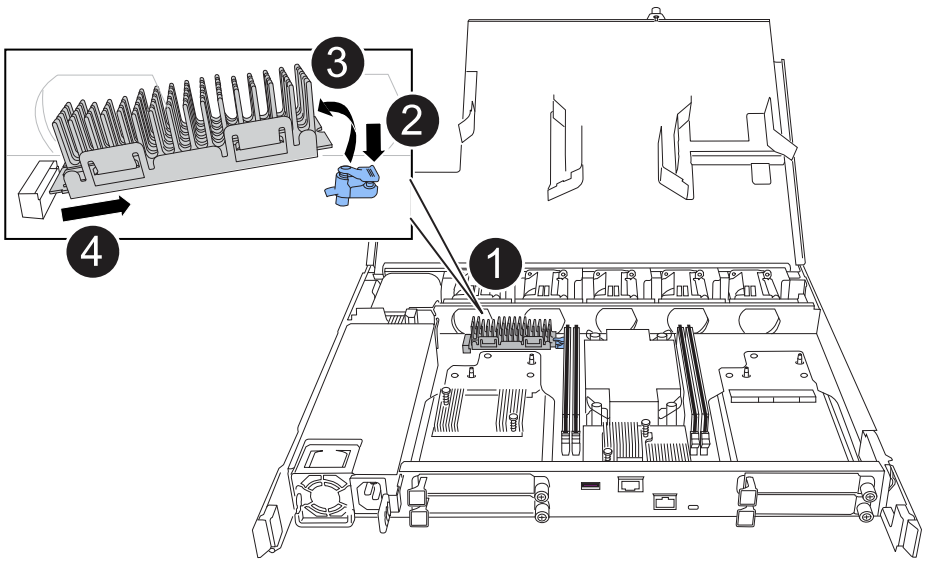

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:
   a. Remove the boot media from its package.
   b. Slide the socket end of the boot media into its socket.
   c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis, but do not reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

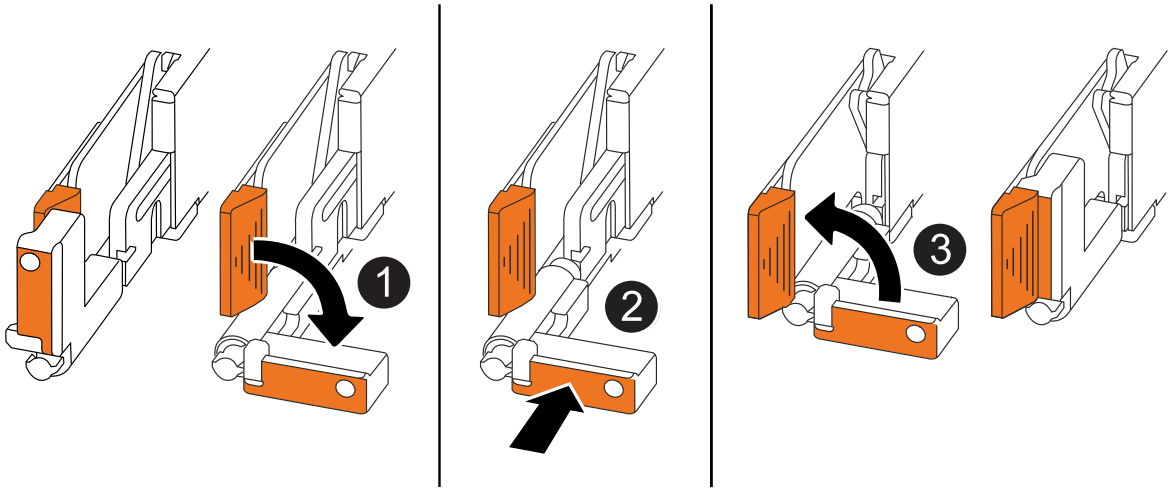

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   > Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.

   > Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

Step 4: Transfer the boot image to the boot media

The replacement boot media that you installed is without an ONTAP image so you need to transfer an ONTAP image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- You must have a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site

- If NVE is supported, download the image with NetApp Volume Encryption, as indicated in the download button.
- If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- You must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image from the Downloads link on the page, to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      The USB flash drive should have the appropriate ONTAP image of what the impaired controller is running.

   c. Remove the USB flash drive from your laptop.
2. Insert the USB flash drive into the USB-A port on the impaired controller.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

3. Fully seat the impaired controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

      The controller boots when fully seated in the chassis. It gets its power from the partner controller.

   b. Rotate the controller handles up and lock in place with the tabs.
4. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

5. Reconnect the power cord to the power supply (PSU) on the impaired controller.

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.

If you are reconnecting a…	Then…
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

After replacing the boot media, you need to boot the recovery image.

Manual boot media recovery from a USB drive - AFF A20, AFF A30, and AFF A50

After installing the new boot media device in your AFF A20, AFF A30, or AFF A50 storage system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `...would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

What's next?

After booting the recovery image, you need to restore encryption on the boot media.

Restore encryption keys after manual boot recovery - AFF A20, AFF A30, and AFF A50

Restore encryption on the replacement boot media in your AFF A20, AFF A30, or AFF A50 storage system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
**************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    `security key-manager key query -restored false.`

    ```
    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    `storage failover giveback -fromnode local`

13. Restore automatic giveback, if you disabled it, by entering the following command.

    `storage failover modify -node local -auto-giveback true`

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    `system node autosupport invoke -node * -type all -message MAINT=END`

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*******************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*******************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

What's next?

After restoring encryption on the boot media, you need to return the failed part to NetApp.

Return the failed part to NetApp - AFF A20, AFF A30, and AFF A50

If a component in your AFF A20, AFF A30, or AFF A50 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - AFF A20, AFF A30, and AFF A50

Get started with replacing the chassis of your AFF A20, AFF A30, or AFF A50 storage system by reviewing the replacement requirements, shutting down the controllers, replacing the chassis, and verifying system operations.

1 Review the chassis replace requirements

To replace the chassis, you must meet certain requirements.

2 Shut down the controllers

Shut down the controllers so you can perform maintenance on the chassis.

3 Replace the chassis

Replacing the chassis includes moving the drives and any drive blanks, controllers (with the power supplies), and bezel from the impaired chassis to the new chassis, and swapping out the impaired chassis with the new chassis of the same model as the impaired chassis.

4 Complete chassis replacement

Verify the HA state of the chassis and return the failed part to NetApp.

Requirements to replace the chassis - AFF A20, AFF A30, and AFF A50

Before replacing the chassis of your AFF A20, AFF A30, or AFF A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement chassis, and the necessary tools.

Review the following requirements and considerations.

Requirements

- The replacement chassis must be the same model as the impaired chassis. This procedure is for a like-for-like replacement, not for an upgrade.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Considerations

- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your storage system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, drives, any drive blanks, and controllers to the new chassis.

What's next?

After you've reviewed the requirements to replace the chassis, you need to shut down the controllers.

Shut down the controllers to replace the chassis - AFF A20, AFF A30, and AFF A50

Shut down the controllers in your AFF A20, AFF A30, or AFF A50 storage system to prevent data loss and ensure system stability when replacing the chassis.

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true
   ```

For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

What's next?

After you've shut down the controllers, you need to replace the chassis.

Replace the chassis - AFF A20, AFF A30, and AFF A50

Replace the chassis of your AFF A20, AFF A30, or AFF A50 storage system when a hardware failure requires it. The replacement process involves removing the controllers, removing the drives, installing the replacement chassis, and reinstalling the chassis components.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

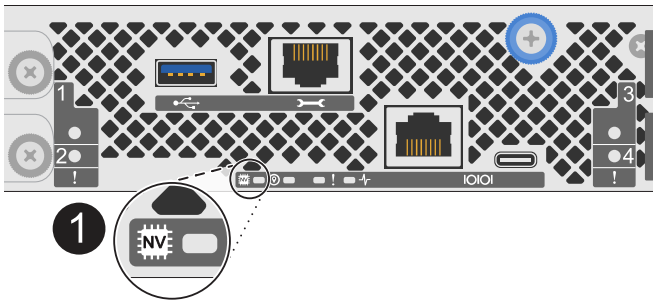

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

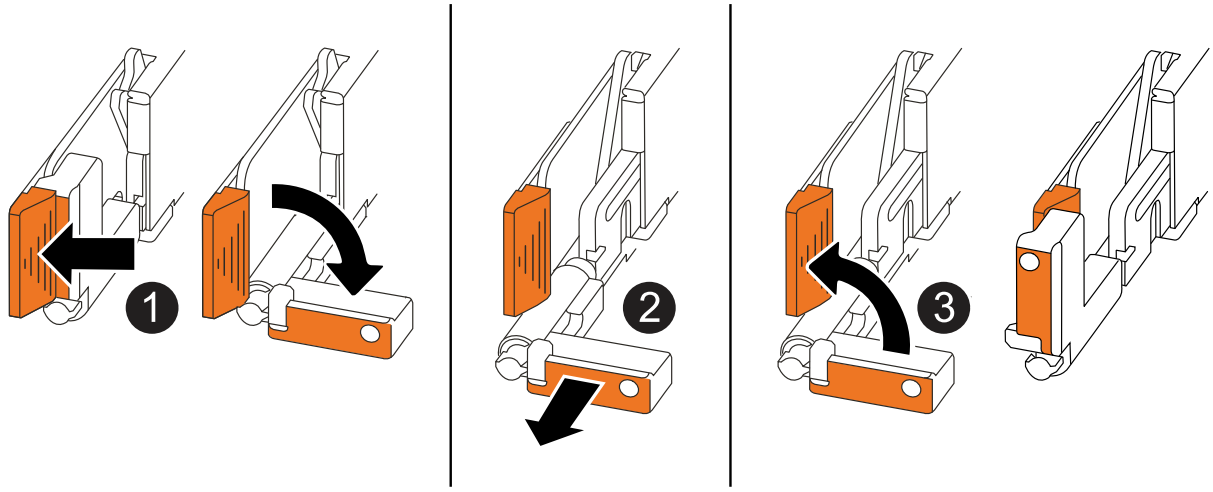

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Repeat these steps for the other controller in the chassis.

Step 2: Remove the drives from the impaired chassis

You need to remove all of the drives and any drive blanks from the impaired chassis so that later in the procedure you can install them in the replacement chassis.

1. Gently remove the bezel from the front of the storage system.
2. Remove the drives and any drive blanks:

Keep track of what drive bay each drive and drive blank was removed from because they must be installed in the same drive bays in the replacement chassis.

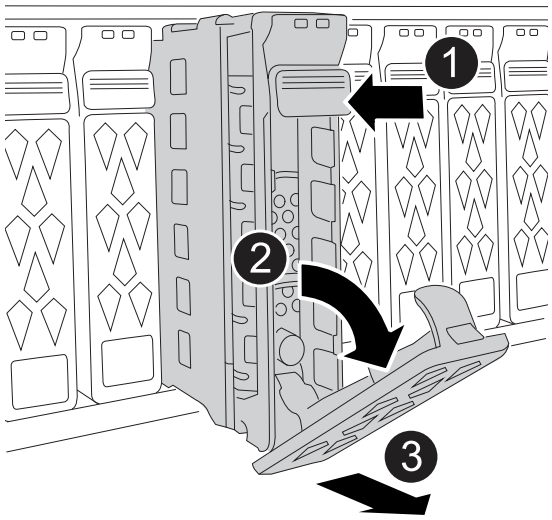

1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.

3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

3. Set the drives aside on a static-free cart or table.

Step 2: Replace the chassis from within the equipment rack or system cabinet

You remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis, install the drives, any drive blanks, and then install the bezel.

1. Remove the screws from the impaired chassis mount points.

   Set the screws aside to use later in this procedure.

   If the storage system shipped in a NetApp system cabinet, you must remove additional screws at the rear of the chassis before the chassis can be removed.

2. Using two people or a power lift, remove the impaired chassis from the equipment rack or system cabinet by sliding it off the rails, and then set it aside.
3. Using two people, install the replacement chassis into the equipment rack or system cabinet by sliding it onto the rails.
4. Secure the front of the replacement chassis to the equipment rack or system cabinet using the screws you removed from the impaired chassis.

Step 4: Install the controllers and drives

Install the controllers and drives into the replacement chassis and reboot the controllers.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when installing a controller, and can be used as a reference for the rest of the controller installation steps.

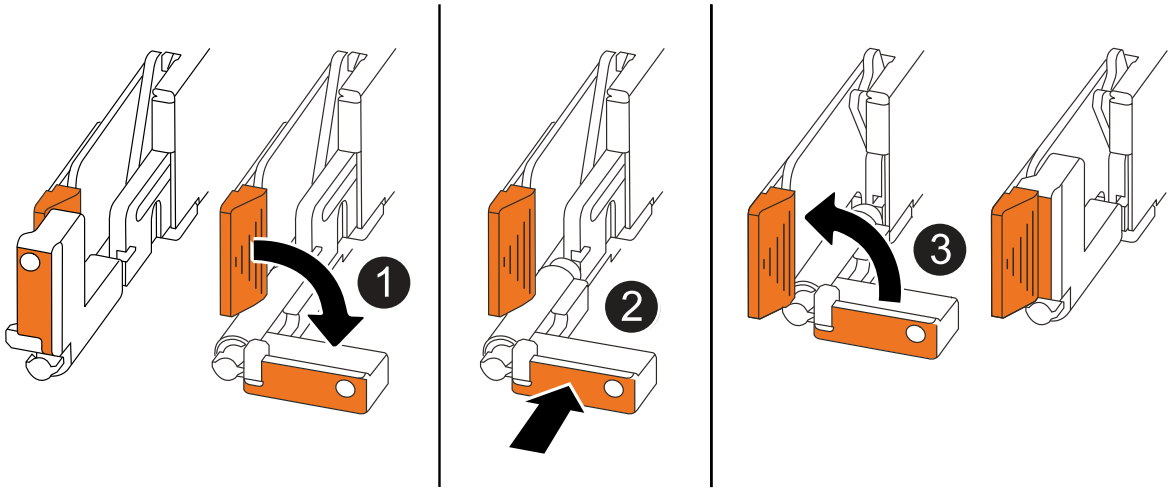

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis and push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

1. Insert one of the controllers into the chassis:
   a. Align the back of the controller with the opening in the chassis.
   b. Firmly push on the handles until the controller meets the midplane and is fully seated in the chassis.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   c. Rotate the controller handles up and lock in place with the tabs.
2. Recable the controller, as needed, except for the power cords.
3. Repeat these steps to install the second controller into the chassis.
4. Install the drives and any drive blanks you removed from the impaired chassis into the replacement chassis:

   The drives and drive blanks must be installed in the same drive bays in the replacement chassis.

   a. With the cam handle in the open position, use both hands to insert the drive.
   b. Gently push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

   d. Repeat the process for the remaining drives.

5. Install the bezel.
6. Reconnect the power cords to the power supplies (PSU) in the controllers.

   Once power is restored to a PSU, the status LED should be green.

   The controllers begin to boot as soon as the power is restored.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. If controllers boot to the LOADER prompt, reboot the controllers:

   ```
   boot_ontap
   ```

8. Turn AutoSupport back on:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

What's next?

After you've replaced the impaired chassis and reinstalled the components into it, you need to complete the chassis replacement.

Complete chassis replacement - AFF A20, AFF A30, and AFF A50

Verify the HA state of the chassis and then return the failed part to NetApp to complete the final step in the AFF A20, AFF A30, and AFF A50 chassis replacement procedure.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your storage system configuration.

1. In Maintenance mode, from either controller, display the HA state of the local controller and chassis:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your storage system configuration:
   a. Set the HA state for the chassis:

      ```
      ha-config modify chassis HA-state
      ```

      The value for HA-state should be *ha*.

The value for HA-state can be one of the following:
* ***ha***
* *mcc* (not supported in ASA)

b. Confirm that the setting has changed:

```
ha-config show
```

3. If you have not already done so, recable the rest of your storage system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - AFF A20, AFF A30, and AFF A50

Get started with replacing the controller in your AFF A20, AFF A30, or AFF A50 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

Review the controller replacement requirements

To replace the controller, you must meet certain requirements.

Shut down the impaired controller

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

Replace the controller

Replacing the controller includes removing the impaired controller, moving FRU components to the replacement controller, installing the replacement controller in the chassis, setting the time and date, and then recabling.

Restore and verify the system configuration

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Give back the controller

Transfer the ownership of storage resources back to the replacement controller.

6 Complete controller replacement

Verify the LIFs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - AFF A20, AFF A30, and AFF A50

Before replacing the controller in your AFF A20, AFF A30, or AFF A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements and considerations for the controller replacement procedure.

Requirements

- All shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.
- You must replace a controller with a controller of the same model type. You cannot upgrade your system by just replacing the controller.
- You cannot change any drives or shelves as part of this procedure.
- You must always capture the controller's console output to a text log file.

  The console output provides you with a record of the procedure you can use to troubleshoot issues you might encounter during the replacement process.

Considerations

It is important that you apply the commands in this procedure to the correct controller:

- The *impaired* controller is the controller that is being replaced.
- The *replacement* controller is the new controller that is replacing the impaired controller.
- The *healthy* controller is the surviving controller.

What's next?

After you've the reviewed the requirements to replace the impaired controller, you need to shut down the impaired controller.

Shut down the impaired controller - AFF A20, AFF A30, and AFF A50

Shut down the impaired controller in your AFF A20, AFF A30, or AFF A50 storage system to prevent data loss and ensure system stability when replacing the controller.

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After you've shut down the impaired controller, you need to replace the controller.

Replace the controller - AFF A20, AFF A30, and AFF A50

Replace the controller in your AFF A20, AFF A30, or AFF A50 storage system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.

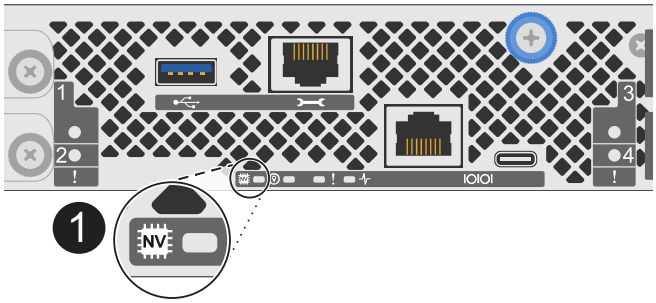

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

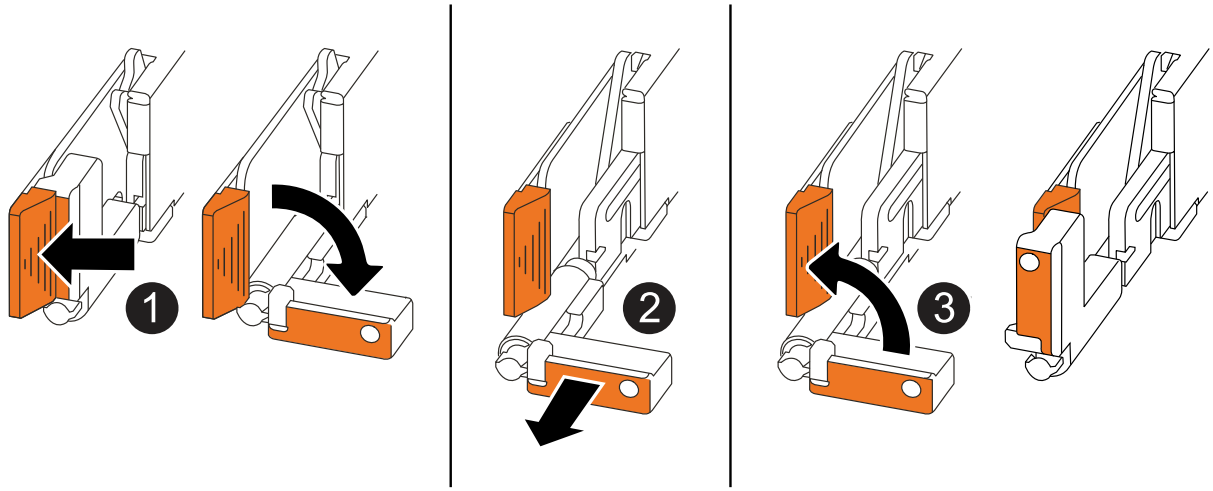

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Move the power supply

Move the power supply (PSU) to the replacement controller.

1. Move the PSU from the impaired controller:

   Make sure the left side controller handle is in the upright position to allow you access to the PSU.

Option 1: Move an AC PSU

To move an AC PSU, complete the following steps.

Steps

1. Remove the AC PSU from the impaired controller:

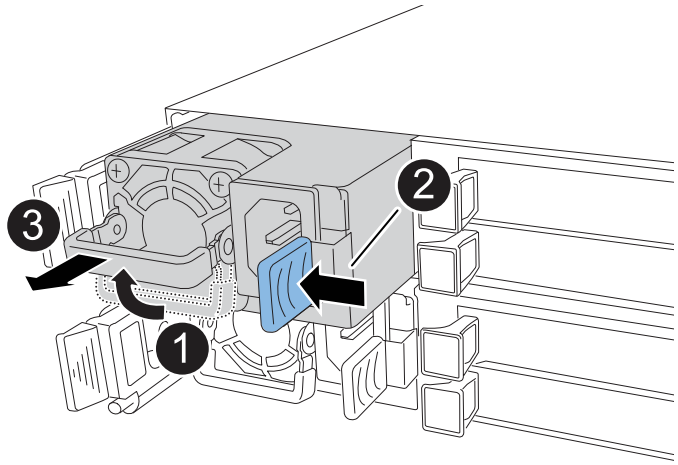

1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

2. Insert the PSU into the replacement controller:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.

Option 2: Move a DC PSU

To move a DC PSU, complete the following steps.

Steps

1. Remove the DC PSU from the impaired controller:

a. Rotate the handle up, to its horizontal position, and then grasp it.

b. With your thumb, press the terracotta tab to release the locking mechanism.

c. Pull the PSU out of the controller while using your other hand to support its weight.

> The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.

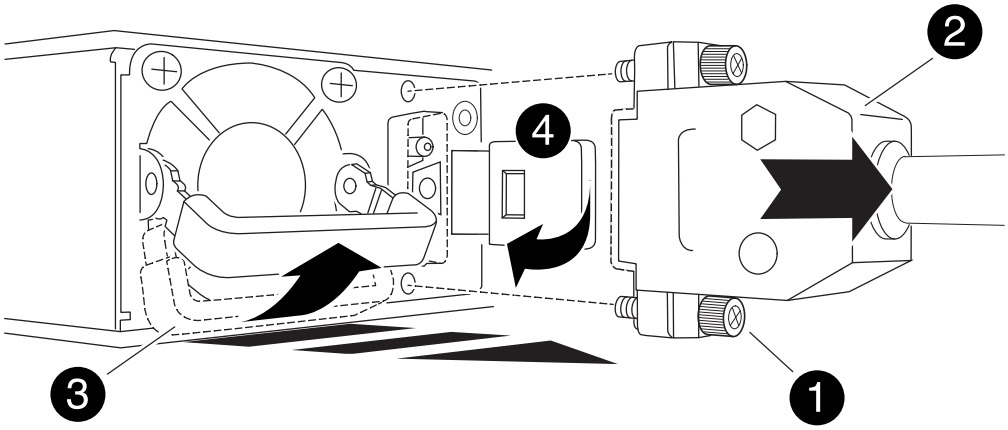

1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

2. Insert the PSU into the replacement controller:

a. Using both hands, support and align the edges of the PSU with the opening in the controller.

b. Gently slide the PSU into the controller until the locking tab clicks into place.

A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

> To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

c. Rotate the handle down, so it is out of the way of normal operations.

Step 3: Move the fans

Move the fans to the replacement controller.

1. Remove one of the fans from the impaired controller:

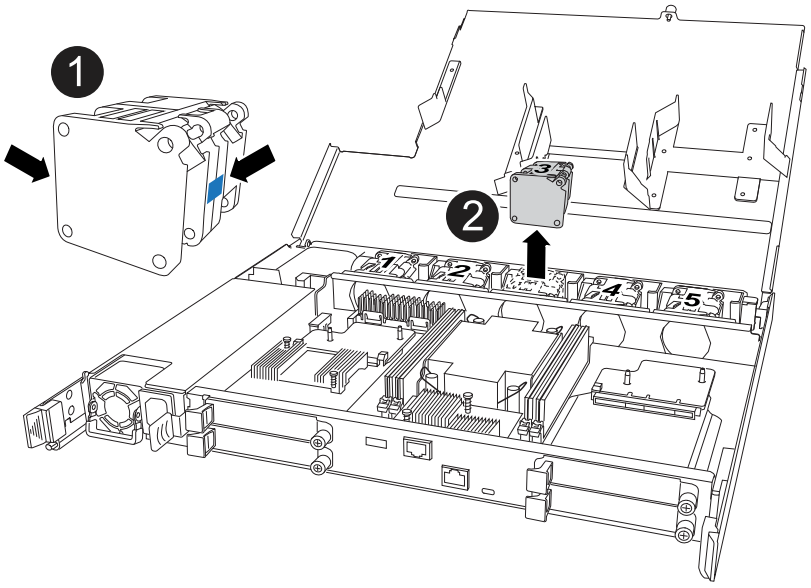

1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

2. Insert the fan into the replacement controller by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.
3. Repeat these steps for the remaining fans.

Step 4: Move the NV battery

Move the NV battery to the replacement controller.

1. Remove the NV battery from the impaired controller:

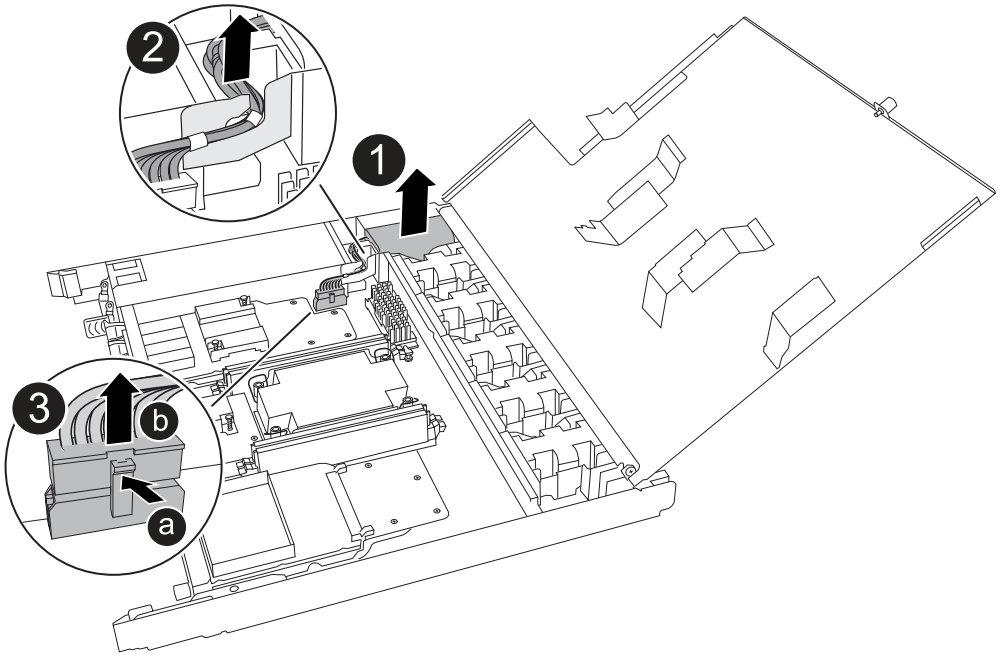

1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

2. Install the NV battery into the replacement controller:
   a. Plug the wiring connector into its socket.
   b. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   c. Place the NV battery into the compartment.

      The NV battery should sit flush in its compartment.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller.

If you have DIMM blanks, you do not need to move them, the replacement controller should come with them installed.

1. Remove one of the DIMMs from the impaired controller:

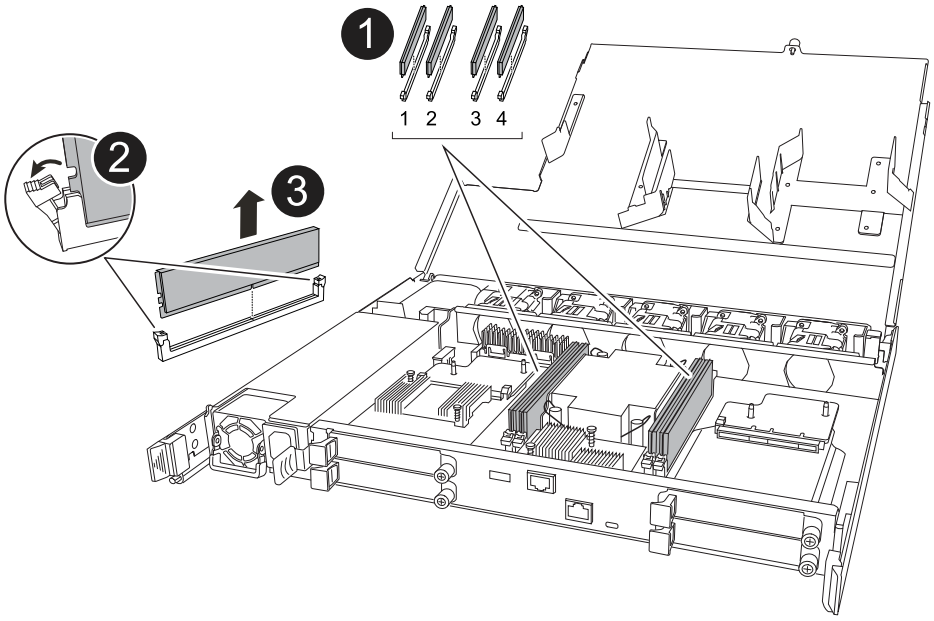

1	DIMM slot numbering and positions. Depending on your storage system model, you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller in the proper orientation. • Eject the DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

2. Install the DIMM in the replacement controller:
   a. Make sure that the DIMM ejector tabs on the connector are in the open position.
   b. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

      When inserted correctly, the DIMM goes in easily but fits tightly in the slot. If not, reinsert the DIMM.

   c. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
   d. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.
3. Repeat these steps for the remaining DIMMs.

Step 6: Move the boot media

Move the boot media to the replacement controller.

1. Remove the boot media from the impaired controller:

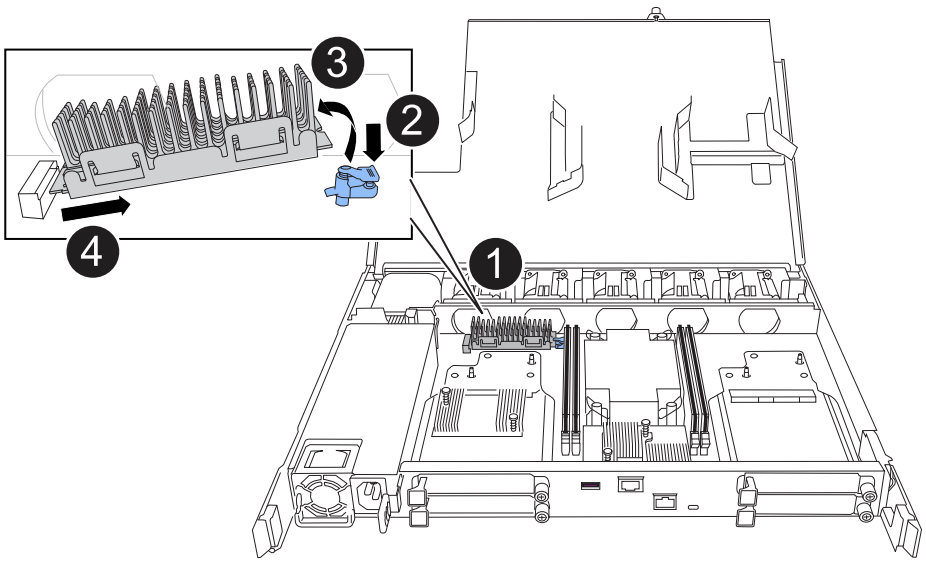

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

2. Install the boot media into the replacement controller:
   a. Slide the socket end of the boot media into its socket.
   b. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 7: Move the I/O modules

Move the I/O modules and any I/O blanking modules to the replacement controller.

1. Unplug cabling from one of the I/O modules.

   Make sure to label the cables so that you know where they came from.

2. Remove the I/O module from the impaired controller:

   Make sure that you keep track of which slot the I/O module was in.

   If you are removing the I/O module in slot 4, make sure the right side controller handle is in the upright position to allow you access to the I/O module.

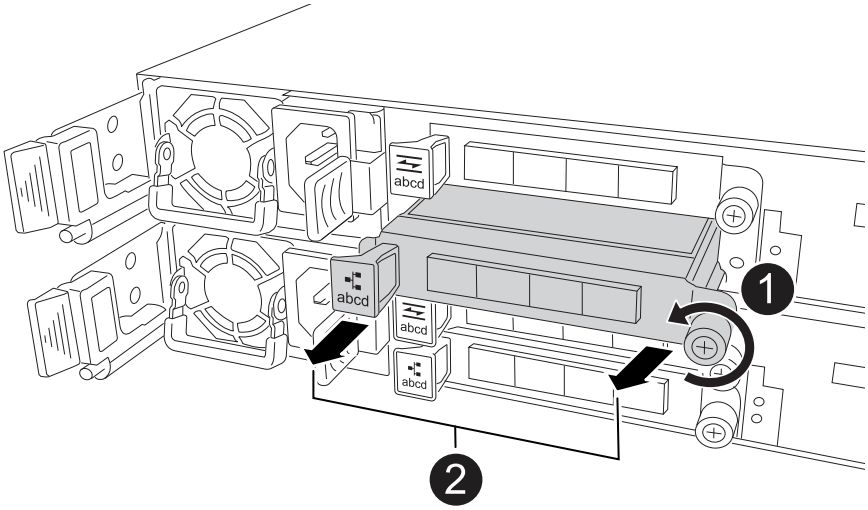

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

3. Install the I/O module into the replacement controller:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
4. Repeat these steps to move the remaining I/O modules and any I/O blanking modules to the replacement controller.

Step 8: Install the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

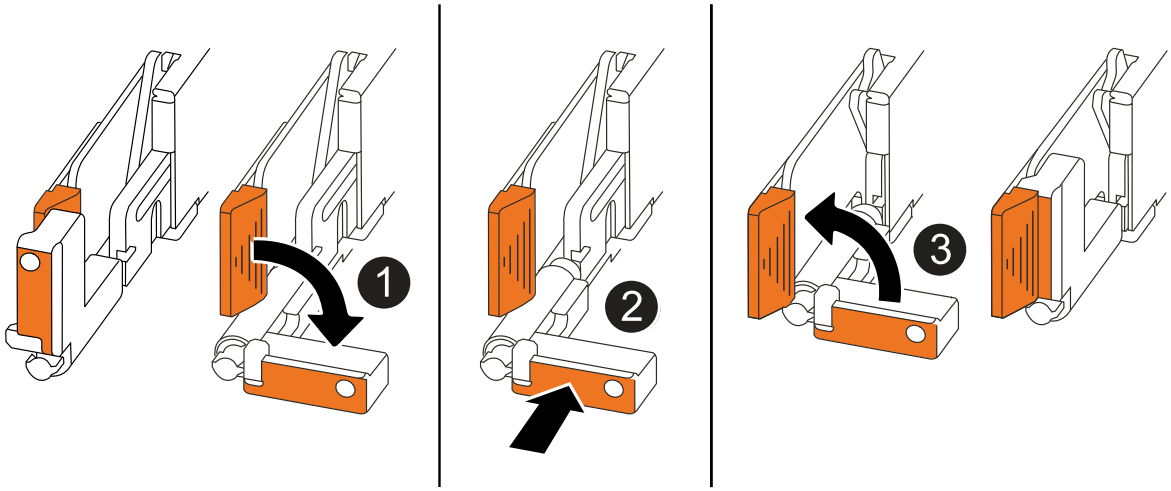

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Take the controller to the LOADER prompt by pressing CTRL-C to abort AUTOBOOT.
6. Set the time and date on the controller:

   Make sure you are at the controller's LOADER prompt.

   a. Display the date and time on the controller:

      ```
      show date
      ```

      Time and date default is in GMT. You have the option to display in local time and in 24hr mode.

   b. Set the current time in GMT:

      ```
      set time hh:mm:ss
      ```

      You can get the current GMT from the healthy node:

      ```
      date -u
      ```

   c. Set the current date in GMT:

      ```
      set date mm/dd/yyyy
      ```

      You can get the current GMT from the healthy node:
      ```
      date -u
      ```

7. Recable the controller as needed.
8. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

After you've replaced the impaired controller, you need to restore the system configuration.

Restore and verify the system configuration - AFF A20, AFF A30, and AFF A50

Verify that the controller's HA configuration is active and functioning correctly in your AFF A20, AFF A30, or AFF A50 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

You must verify the `HA` state of the controller and, if necessary, update the state to match your storage system configuration.

1. Boot to maintenance mode:

   ```
   boot_ontap maint
   ```

   a. Enter `y` when you see *Continue with boot?*.

      If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

5. If the displayed system state of the controller does not match your storage system configuration, set the `HA` state for the controller:

   ```
   ha-config modify controller ha
   ```

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed:

   ```
   ha-config show
   ```

Step 2: Verify disk list

1. Verify that the adapter lists the paths to all disks:

   ```
   storage show disk -p
   ```

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode:

```
halt
```

What's next?

After you've restored and verified your system configuration, you need to give back the controller.

Give back the controller - AFF A20, AFF A30, and AFF A50

Return control of storage resources to the replacement controller so your AFF A20, AFF A30, or AFF A50 storage system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption, Onboard Key Manager (OKM) encryption, or External Key Manager (EKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`

   If you encounter errors, contact NetApp Support.

9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

c. Verify status of the keys: `security key-manager key query -restored *false*`

   The output should show no results when when properly synchronized.

   If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode *impaired_node_name*`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback *true*`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

External key manager (EKM)

Reset encryption and return the controller to normal operation.

Steps

1. If the root volume is encrypted with External Key Manager and the console cable is connected to the replacement node, enter `boot_ontap menu` and select option `11`.
2. If these questions appear, answer `y` or `n` as appropriate:

   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}

   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}

   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}

   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}

   Do you know the KMIP server address? {y/n}

   Do you know the KMIP port? {y/n}

   Contact NetApp Support if you have issues.

3. Supply the information for:
   - The client certificate (client.crt) file contents
   - The client key (client.key) file contents
   - The KMIP server CA(s) (CA.pem) file contents
   - The IP address for the KMIP server
   - The port for the KMIP server
4. Once the system processes, you see the Boot Menu. Select '1' for normal boot.
5. Check the takeover status: `storage failover show`
6. Ensure any core dumps on the repaired node are saved by going to advanced mode `set -privilege advanced` and then `run local partner nosavecore`.
7. Return the impaired controller to normal operation by giving back its storage: `storage failover`

```
giveback -ofnode impaired_node_name
```

8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
9. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

After you've transferred the ownership of storage resources to the replacement controller, you need to complete the controller replacement procedure.

Complete controller replacement - AFF A20, AFF A30, and AFF A50

To complete the controller replacement for your AFF A20, AFF A30, or AFF A50 storage system, first restore the NetApp Storage Encryption configuration (if necessary) and install the required licenses on the new controller. Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, register the new controller's serial number and then return the failed part to NetApp.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

  Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The licenses keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.
- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs, register the serial number, and check cluster health

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A20, AFF A30, and AFF A50

Replace a DIMM in your AFF A20, AFF A30, or AFF A50 storage system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- All other components in the storage system must be working correctly; if not, contact NetApp Support before continuing.

- You must replace the failed FRU component with a replacement FRU component you received from your provider.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

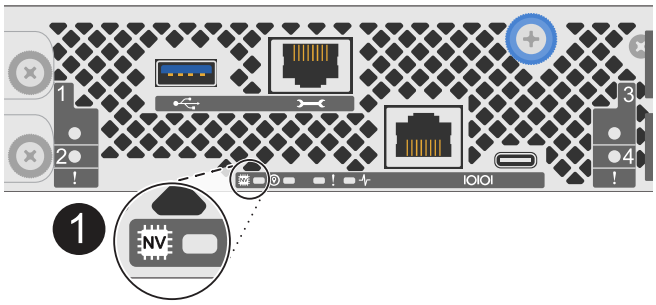

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

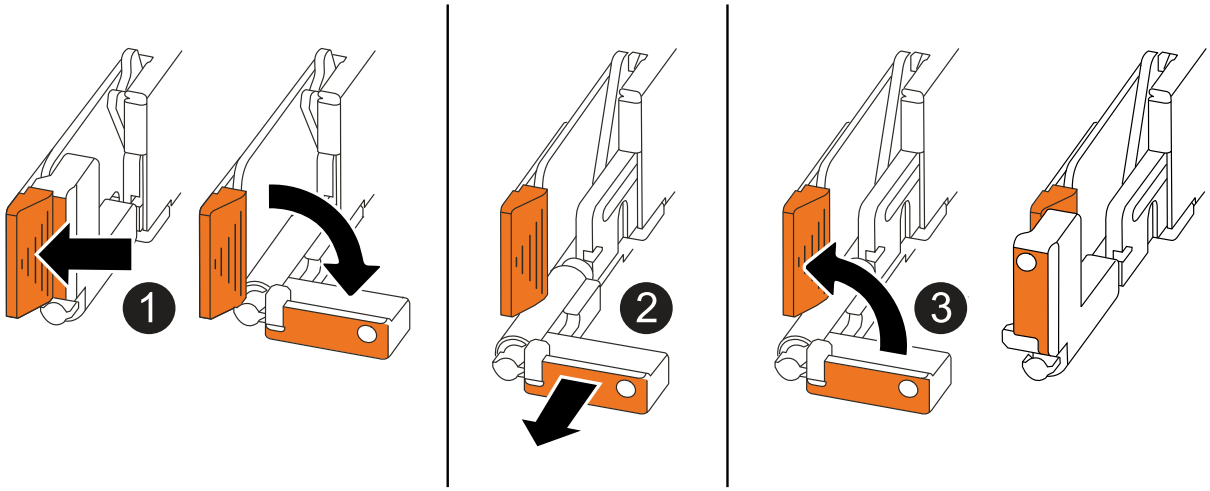

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace a DIMM

To replace a DIMM, locate the faulty DIMM inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller and identify the faulty DIMM.

> Consult either the Netapp Hardware Universe or the FRU map on the cover of the controller for exact DIMM locations.

3. Remove the faulty DIMM:

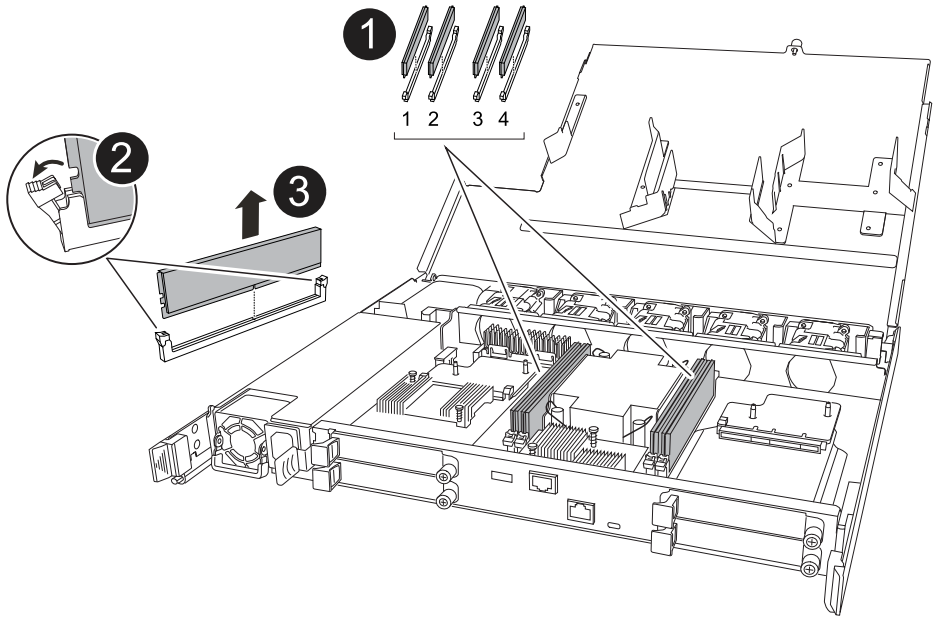

1	DIMM slot numbering and positions. Depending on your storage system model you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM using the same orientation. • Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

4. Install the replacement DIMM:
   a. Remove the replacement DIMM from its antistatic shipping bag.
   b. Make sure that the DIMM ejector tabs on the connector are in the open position.
   c. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

      When inserted correctly, the DIMM goes in easily but fits tightly in the slot. Reinsert the DIMM if you feel it is not inserted correctly.

   d. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
   e. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

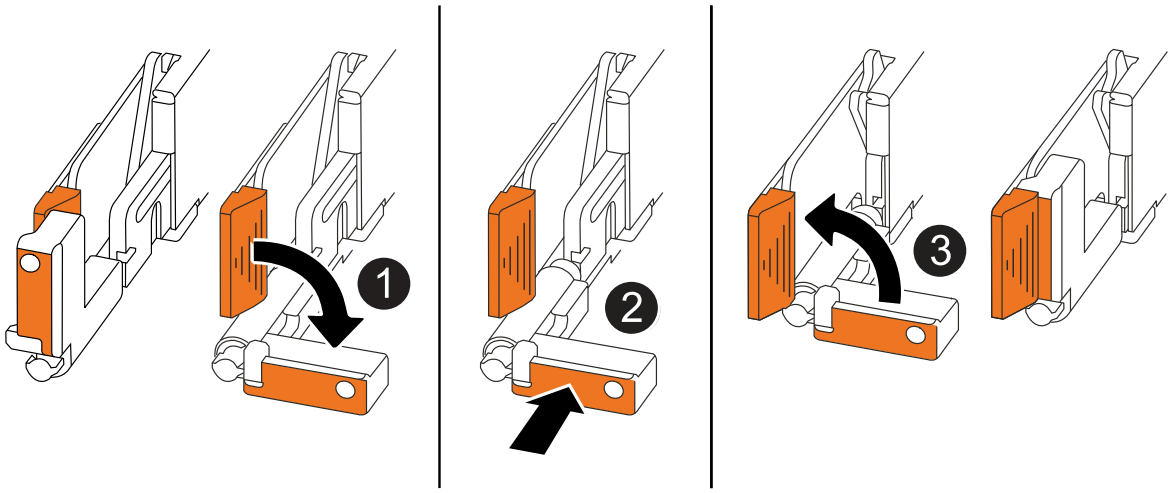

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

> Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

> Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

> Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

> The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a Drive - AFF A20, AFF A30, and AFF A50

Replace a drive in your AFF A20, AFF A30, or AFF A50 storage system when a drive fails or requires an upgrade. The replacement process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed SSD drive nondisruptively while I/O is in progress.

Before you begin

- The drive that you are installing must be supported by your storage system.

  NetApp Hardware Universe

- If self-encrypting drive (SED) authentication is enabled, you must use the SED replacement instructions in the ONTAP documentation.

  Instructions in the ONTAP documentation describe additional steps you must perform before and after replacing an SED.

  NetApp encryption overview with the CLI

- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.
- Verify that the drive you are removing is failed.

  You can verify that the drive is failed by running the `storage disk show -broken` command. The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  Depending on the drive type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

About this task

- When replacing a failed drive, you must wait 70 seconds between the removal of the drive and the insertion of the replacement drive to allow the storage system to recognize that a drive was removed.
- The best practice is to have the current version of the Disk Qualification Package (DQP) installed before hot-swapping a drive.

  Having the current version of the DQP installed allows your system to recognize and use newly qualified drives. This avoids system event messages about having noncurrent drive information and prevention of drive partitioning because drives are not recognized. The DQP also notifies you of noncurrent drive firmware.

  NetApp Downloads: Disk Qualification Package

- The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on

your system before replacing FRU components.

NetApp Downloads: Disk Shelf Firmware

NetApp Downloads: Disk Drive Firmware

> Do not revert firmware to a version that does not support your shelf and its components.

- Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

  > Drive firmware checks occur every two minutes.

- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment if it is enabled.

   > You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled:

      ```
      storage disk option show
      ```

      You can enter the command on either controller.

      If automatic drive assignment is enabled, the output shows `on` in the `Auto Assign` column (for each controller).

   b. If automatic drive assignment is enabled, disable it:

      ```
      storage disk option modify -node node_name -autoassign off
      ```

      You must disable automatic drive assignment on both controllers.

2. Properly ground yourself.
3. Remove the bezel from the front of the storage system.
4. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

5. Remove the failed drive:

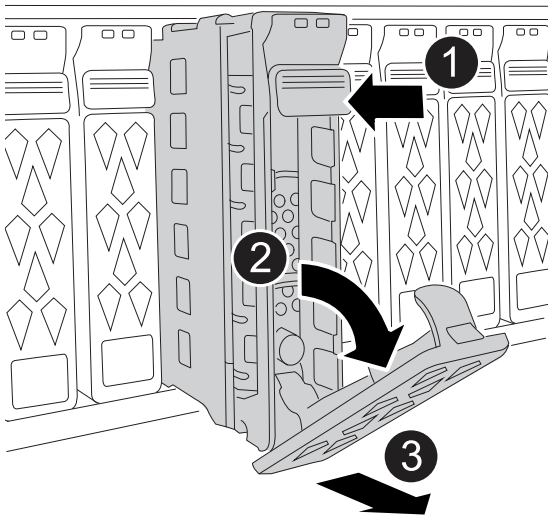

1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

6. Wait a minimum of 70 seconds before inserting the replacement drive.
7. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the drive.
   b. Gently push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
8. Verify that the drive's activity (green) LED is illuminated.

   When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

9. If you are replacing another drive, repeat steps 4 through step 8.
10. Reinstall the bezel on the front of the storage system.
11. If you disabled automatic drive assignment in step 1, manually assign drive ownership, and then reenable automatic drive assignment if needed:
    a. Display all unowned drives:

       ```
       storage disk show -container-type unassigned
       ```

       You can enter the command on either controller.

    b. Assign each drive:

       ```
       storage disk assign -disk disk_name -owner owner_name
       ```

       You can enter the command on either controller.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed:

       ```
       storage disk option modify -node node_name -autoassign on
       ```

       You must reenable automatic drive assignment on both controllers.
12. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan module - AFF A20, AFF A30, and AFF A50

Replace a fan module in your AFF A20, AFF A30, or AFF A50 storage system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

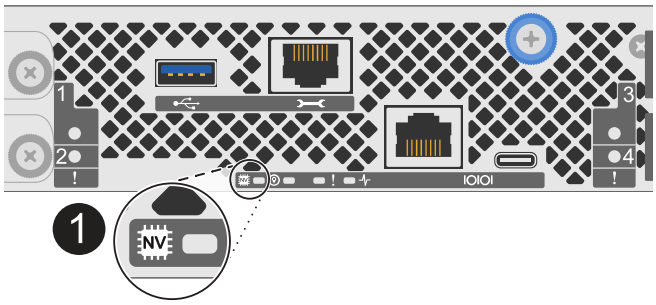

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

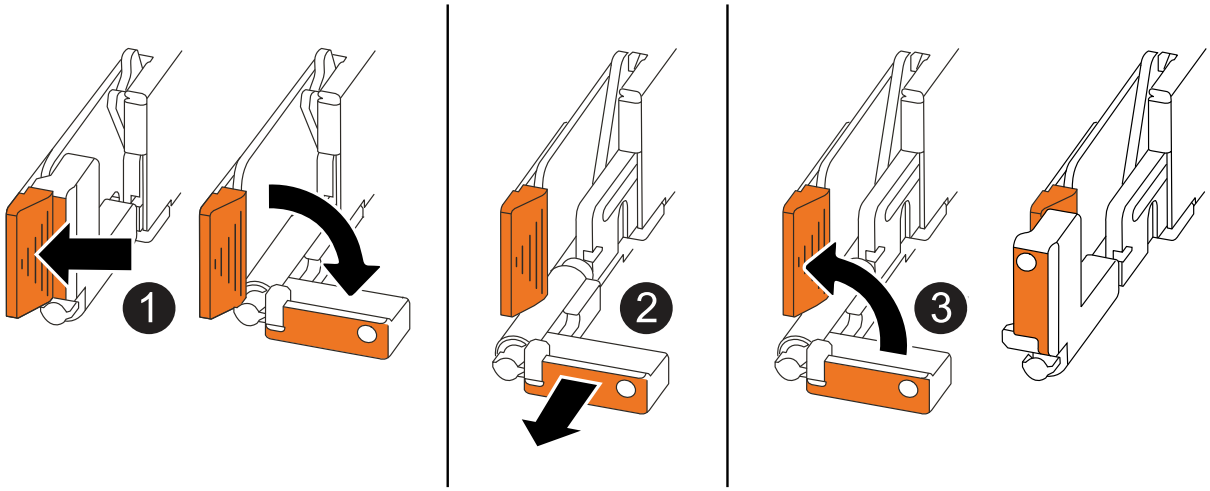

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace fan

To replace a fan, remove the failed fan and replace it with a new fan.

Steps

1. Identify the fan that you must replace by checking the console error messages.
2. Remove the failed fan:

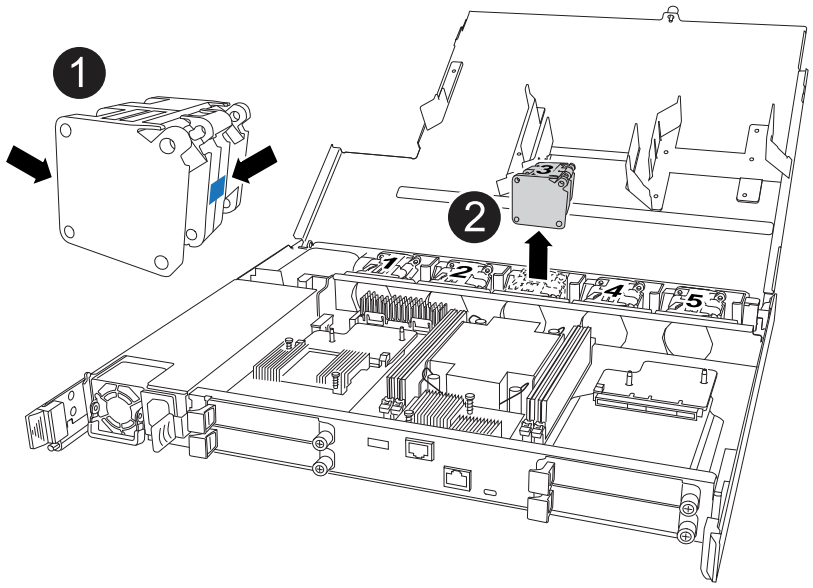

1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

3. Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

Step 4: Reinstall the controller module

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

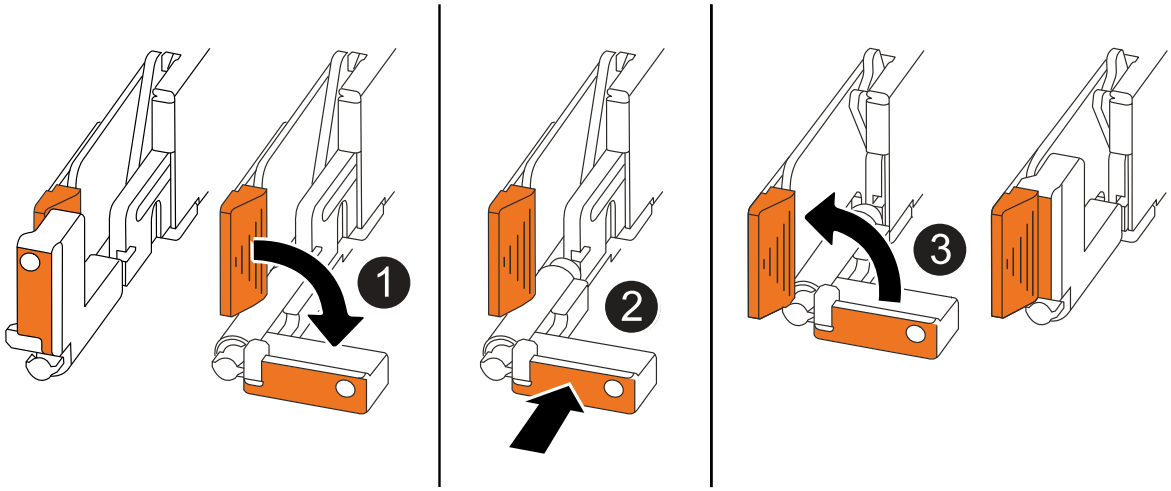

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Overview of I/O module maintenance - AFF A20, AFF A30, and AFF A50

The AFF A20, AFF A30, and AFF A50 storage systems offer flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding, hot-swapping, or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your storage system with the same type of I/O module, or with a different type of I/O module. You can hot-swap a cluster and HA I/O module when your storage system meets specific requirements. You can also add an I/O module to a storage system with available slots.

- Add an I/O module

  Adding additional I/O modules can improve redundancy, helping to ensure that the storage system remains operational even if one I/O module fails.

- Hot-swap a cluster and HA I/O module

Hot-swapping a failed cluster and HA I/O module can restore the storage system to its optimal operating state. Hot-swapping is done without having to manually take over the impaired controller.

To use this procedure, your storage system must be running ONTAP 9.17.1 or later and meet specific storage system requirements.

- Replace an I/O module

  Replacing a failing I/O module can restore the storage system to its optimal operating state.

Add an I/O module - AFF A20, AFF A30, and AFF A50

Add an I/O module to your AFF A20, AFF A30, or AFF A50 storage system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your AFF A20, AFF A30, and AFF A50 storage systems when there are slots available or when all slots are fully populated.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller module

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, remove the I/O blanking module from the target slot.

   Unused I/O slots should have blanking module installed to prevent possible thermal issues and for EMC compliance.

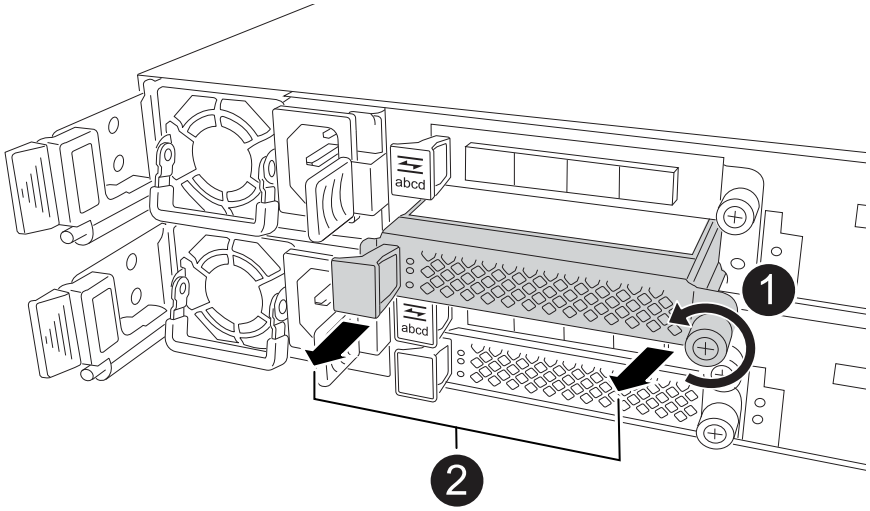

1	On the I/O blanking module, turn the thumbscrew counterclockwise to loosen.
2	Pull the I/O blanking module out of the controller using the tab on the left and the thumbscrew.

3. Install the new I/O module:
   a. Align the I/O module with the edges of the controller slot opening.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O module.

   c. Turn the thumbscrew clockwise to tighten.
4. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.

5. Reboot the impaired controller from the LOADER prompt: `bye`

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

6. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name.
```

7. Repeat these steps to add an I/O module to the other controller.
8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation: +

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, unplug any cabling on the target I/O module.
3. Remove the target I/O module from the controller:

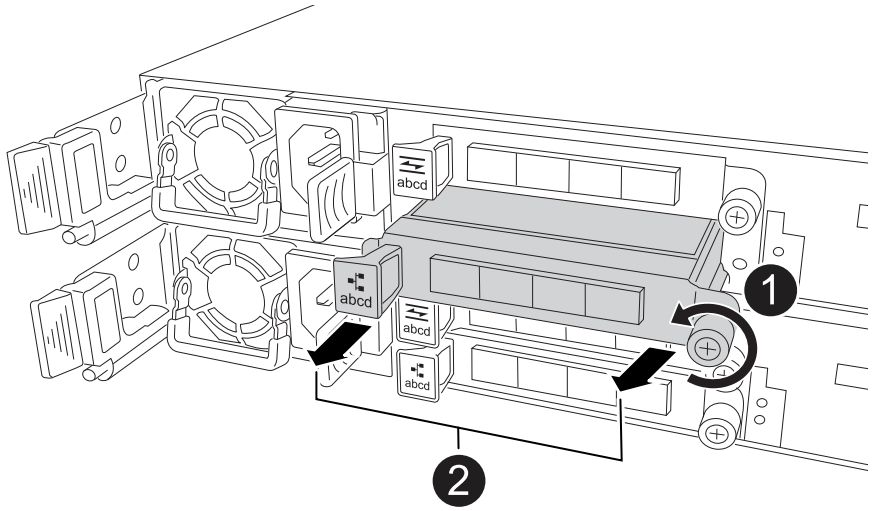

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the new I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
5. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
6. Repeat the I/O module remove and install steps to add any additional I/O modules in the controller.
7. Reboot the impaired controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.
8. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```
9. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```
10. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```
11. If you installed a NIC module, specify the usage mode for each port as *network*:

    ```
    storage port modify -node node_name -port port_name -mode network
    ```
12. Repeat these steps for the other controller.

Hot-swap an I/O module used for cluster and HA traffic - AFF A20, AFF A30, and AFF A50

The cluster and HA I/O module supports interconnects for clustering and high-availability. You can hot-swap the module in your AFF A20, AFF A30, or AFF A50 storage system when the module fails and if your storage system meets specific requirements.

To hot-swap a module, you ensure your storage system meets the procedure requirements, prepare the storage system and I/O module in slot 4, hot-swap the failed module for an equivalent one, bring the replacement module online, restore the storage system to normal operation, and return the failed module to NetApp.

About this task

- Hot-swapping the cluster and HA I/O module means that you do not have to perform a manual takeover; the impaired controller (the controller with the failed cluster and HA I/O module) has automatically taken over the healthy controller.

  When the impaired controller has taken over the healthy controller, the only way to recover without an outage is to hot-swap the module.

- It is critical to apply the commands to the correct controller when you are hot-swapping the cluster and HA I/O module:
  - The *impaired controller* is the controller on which you are hot-swapping the cluster and HA I/O module and it is the controller that has taken over the healthy controller.
  - The *healthy controller* is the HA partner of the impaired controller and it is the controller that was taken over by the impaired controller.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, make sure your storage system meets all requirements.

> If your storage system does not meet all requirements, you must use the replace an I/O module procedure.

- Your storage system must be running ONTAP 9.17.1 or later.
- The I/O module that failed must be a cluster and HA I/O module in slot 4 and you must be replacing it with an equivalent cluster and HA I/O module. You cannot change the I/O module type.
- Your storage system configuration must have only one cluster and HA I/O module located in slot 4, not two cluster and HA I/O modules.
- Your storage system must be a two-node (switchless or switched) cluster configuration.
- The controller with the failed cluster and HA I/O module (the impaired controller) must have already taken over the healthy partner controller. The takeover should have occurred automatically if the I/O module is failed.

  For two-node clusters, the storage system cannot discern which controller has the failed I/O module, so either controller might initiate the takeover. The cluster and HA I/O module hot-swap procedure is only supported when the controller with the failed I/O module (the impaired controller) has taken over the healthy controller.

You can verify that the impaired controller successfully took over the healthy controller by entering the `storage failover show` command.

If you are not sure which controller has the failed I/O module, contact NetApp Support.

- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Step 2: Prepare the storage system and I/O module slot 4

Prepare the storage system and I/O module slot 4 so that it is safe to remove the failed cluster and HA I/O module:

Steps

1. Properly ground yourself.
2. Unplug cabling from the failed cluster and HA I/O module.

   Make sure to label the cables so that later in this procedure you can reconnect them to the same ports.

3. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<number of hours down>h`

   For example, the following AutoSupport message suppresses automatic case creation for two hours:

   `node2::> system node autosupport invoke -node * -type all -message MAINT=2h`

4. Disable automatic giveback:
   a. Enter the following command from the console of the impaired controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
5. Prepare the failed cluster and HA module in slot 4 for removal by removing it from service and powering it off:
   a. Enter the following command:

      `system controller slot module remove -node impaired_node_name -slot slot_number`

   b. Enter `y` when you see the prompt *Do you want to continue?*

      For example, the following command prepares the module in slot 4 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

```
node2::> system controller slot module remove -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered off for removal.

Do you want to continue? {y|n}: y

The module has been successfully removed from service and powered
off. It can now be safely removed.
```

6. Verify the failed cluster and HA module in slot 4 is powered off:

   `system controller slot module show`

   The output should show *`powered-off`* in the status column for the failed module in slot 4.

Step 3: Replace the failed cluster and HA I/O module

Replace the failed cluster and HA I/O module in slot 4 with an equivalent I/O module:

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the failed cluster and HA I/O module from the impaired controller:

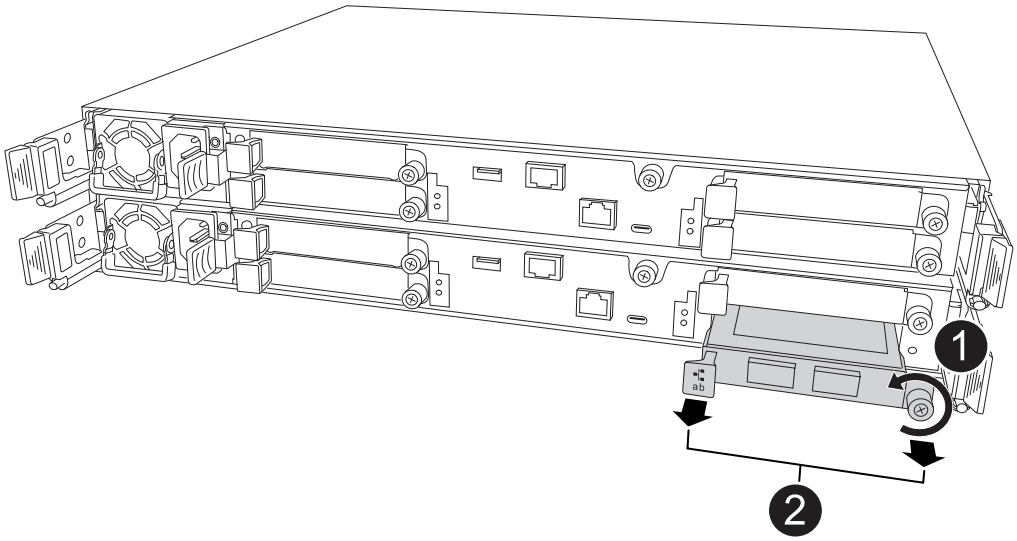

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew on the right.

3. Install the replacement cluster and HA I/O module into slot 4:

a. Align the I/O module with the edges of the slot.
b. Gently push the I/O module all the way into the slot, making sure to properly seat the I/O module into the connector.

   You can use the tab on the left and the thumbscrew on the right to push in the I/O module.

c. Turn the thumbscrew clockwise to tighten.

4. Cable the cluster and HA I/O module.

Step 4: Bring the replacement cluster and HA I/O module online

Bring the replacement cluster and HA I/O module in slot 4 online, verify the module ports initialized successfully, verify slot 4 is powered on, and then verify the module is online and recognized.

Steps

1. Bring the replacement cluster and HA I/O module online:
   a. Enter the following command:

      `system controller slot module insert -node impaired_node_name -slot slot_name`

   b. Enter `y` when you see the prompt, *Do you want to continue?*

      The output should confirm the cluster and HA I/O module was successfully brought online (powered on, initialized, and placed into service).

      For example, the following command brings slot 4 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered on and initialized.

Do you want to continue? {y|n}: `y`

The module has been successfully powered on, initialized and placed
into service.
```

2. Verify that each port on the cluster and HA I/O module successfully initialized:

   `event log show -event *hotplug.init*`

   > It might take several minutes to allow for any required firmware updates and port initialization.

   The output should show a hotplug.init.success EMS event logged for each port on the cluster and HA I/O module with *`hotplug.init.success:`* in the *`Event`* column.

For example, the following output shows initialization succeeded for cluster and HA I/O module ports e4b and e4a:

```
node2::> event log show -event *hotplug.init*

Time                Node             Severity      Event

------------------- ---------------- -------------
---------------------------

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4b" in slot 4 succeeded

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4a" in slot 4 succeeded

2 entries were displayed.
```

3. Verify I/O module slot 4 is powered on and ready for operation:

   `system controller slot module show`

   The output should show slot 4 status as *`powered-on`* and therefore ready for operation of the replacement cluster and HA I/O module.

4. Verify that the replacement cluster and HA I/O module is online and recognized.

   Enter the command from the console of the impaired controller:

   `system controller config show -node local -slot4`

   If the replacement cluster and HA I/O module was successfully brought online and is recognize, the output shows I/O module information, including port information, for slot 4.

   For example, you should see output similar to the following:

```
node2::> system controller config show -node local -slot 4

Node: node2
Sub- Device/
Slot slot Information
---- ---- ------------------------------
   4    - Dual 40G/100G Ethernet Controller CX6-DX
                  e4a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2807GJFM-B
                  e4b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2809G26F-A
                  Device Type:        CX6-DX PSID(NAP0000000027)
                  Firmware Version:   22.44.1700
                  Part Number:        111-05341
                  Hardware Revision:  20
                  Serial Number:      032403001370
```

Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the healthy controller, restoring automatic giveback, and reenabling AutoSupport automatic case creation.

Steps

1. Return the healthy controller (the controller that was taken over) to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode healthy_node_name
   ```

2. Restore automatic giveback from the console of the impaired controller (the controller that took over the healthy controller):

   ```
   storage failover modify -node local -auto-giveback true
   ```

3. If AutoSupport is enabled, restore automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=end
   ```

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an I/O module - AFF A20, AFF A30, and AFF A50

Replace an I/O module in your AFF A20, AFF A30, or AFF A50 storage system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

Before you begin

All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace a failed I/O module

To replace a failed I/O module, locate it in the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug cabling from the failed I/O module.

   Make sure to label the cables so that you know where they came from.

3. Remove the failed I/O module from the controller:

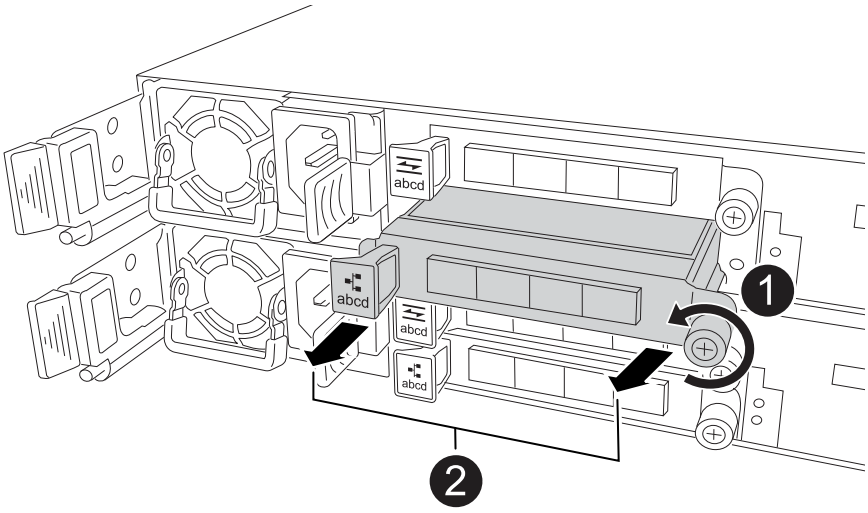

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the replacement I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

c. Turn the thumbscrew clockwise to tighten.

5. Cable the I/O module.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt:

```
bye
```

Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

3. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

4. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NV battery - AFF A20, AFF A30, and AFF A50

Replace the NV battery in your AFF A20, AFF A30, or AFF A50 storage system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

Before you begin

All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

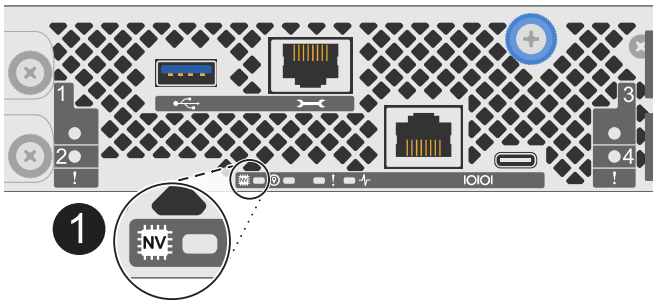

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

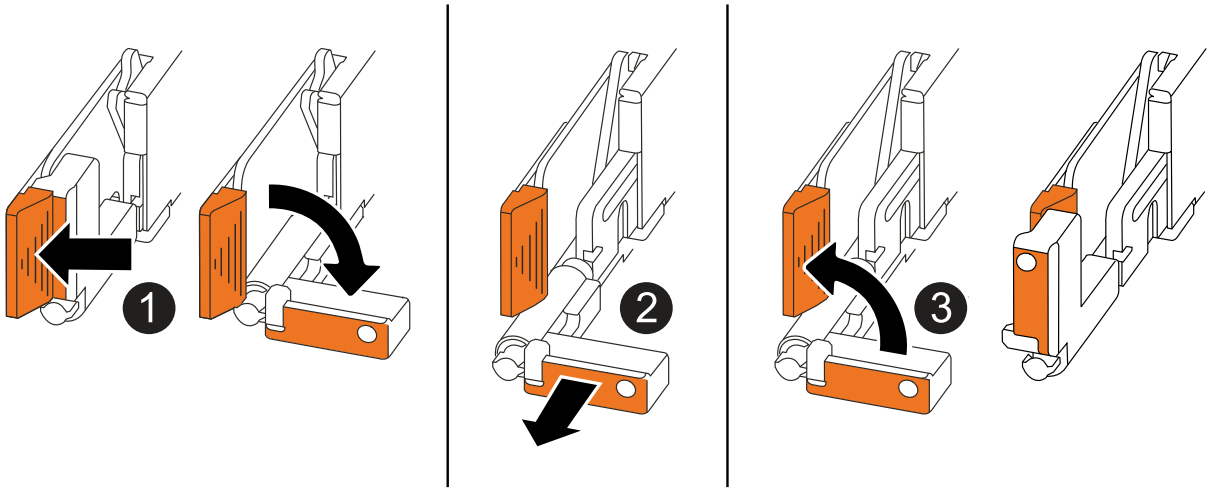

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the NV battery

Remove the failed NV battery from the controller and install the replacement NV battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the NV battery.
3. Remove the NV battery:

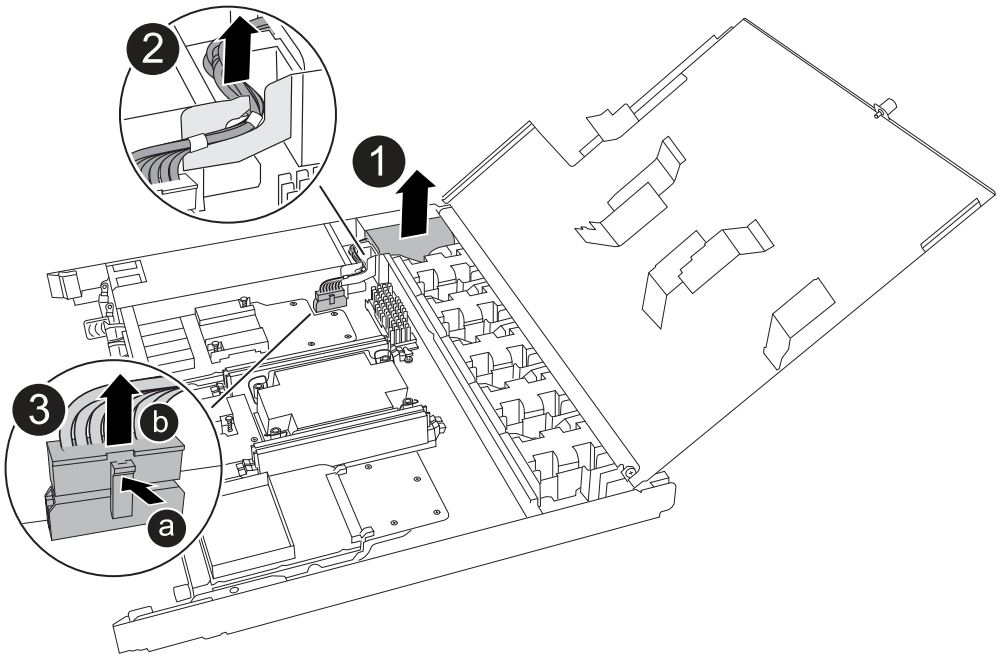

1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

4. Install the replacement NV battery:
   a. Remove the replacement battery from its package.
   b. Plug the wiring connector into its socket.
   c. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   d. Place the NV battery into its compartment.

      The NV battery should sit flush in its compartment.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

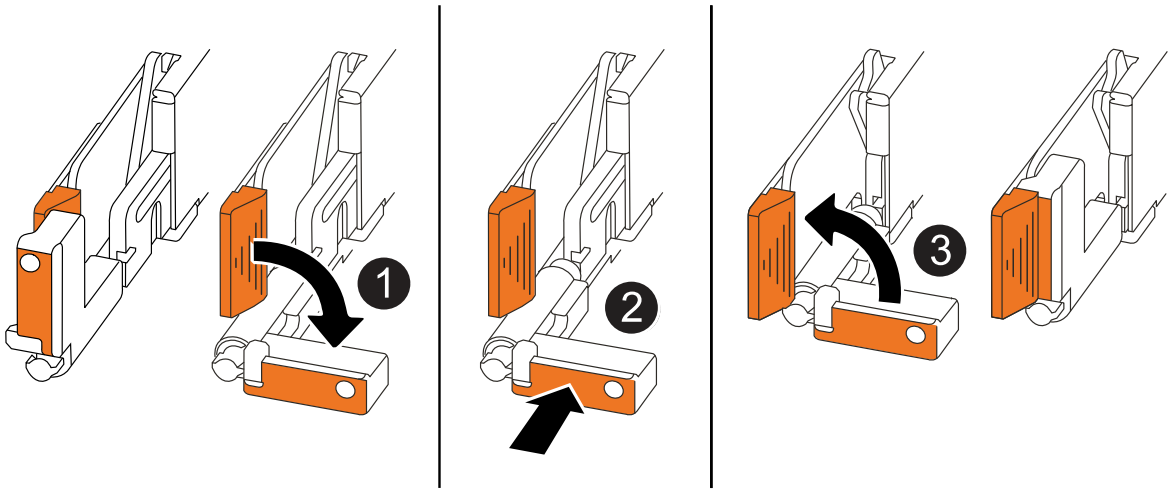

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.

The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - AFF A20, AFF A30, and AFF A50

Replace an AC or DC power supply unit (PSU) in your AFF A20, AFF A30, or AFF A50 storage system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cord, replacing the faulty PSU, and then reconnecting it to the power source.

About this task

- This procedure is written for replacing one PSU at a time.

The PSUs are redundant and hot-swappable.

- IMPORTANT: Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.
- Use the appropriate procedure for your type of PSU: AC or DC.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the power cord from the PSU by opening the power cord retainer, and then unplug the power cord from the PSU.

   PSUs do not have a power switch.

4. Remove the PSU:

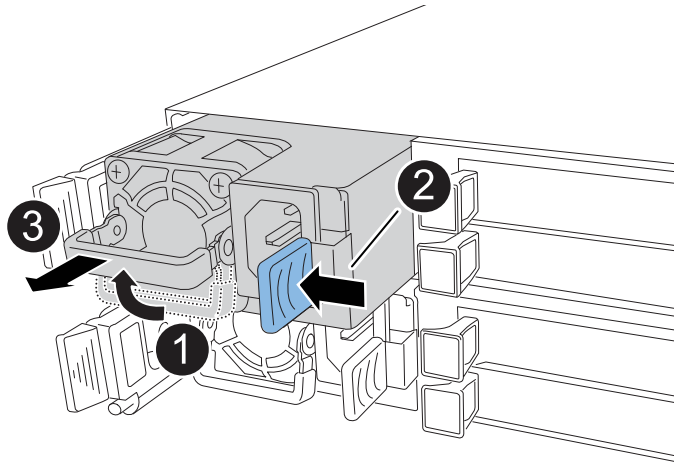

1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

5. Install the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

c. Rotate the handle down, so it is out of the way of normal operations.

6. Reconnect the power cord to the PSU and secure the power cord with the power cord retainer.

   Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the PSU:

   PSUs do not have a power switch.

   a. Unscrew the two thumb screws on the D-SUB DC power cord connector.

      The illustration and table in step 4 shows the two thumb screws (item #1) and the D-SUB DC power cord connector (item #2).

   b. Unplug the cord from the PSU and set it aside.

4. Remove the PSU:

   a. Rotate the handle up, to its horizontal position, and then grasp it.

   b. With your thumb, press the terracotta tab to release the locking mechanism.

   c. Pull the PSU out of the controller while using your other hand to support its weight.

   The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.

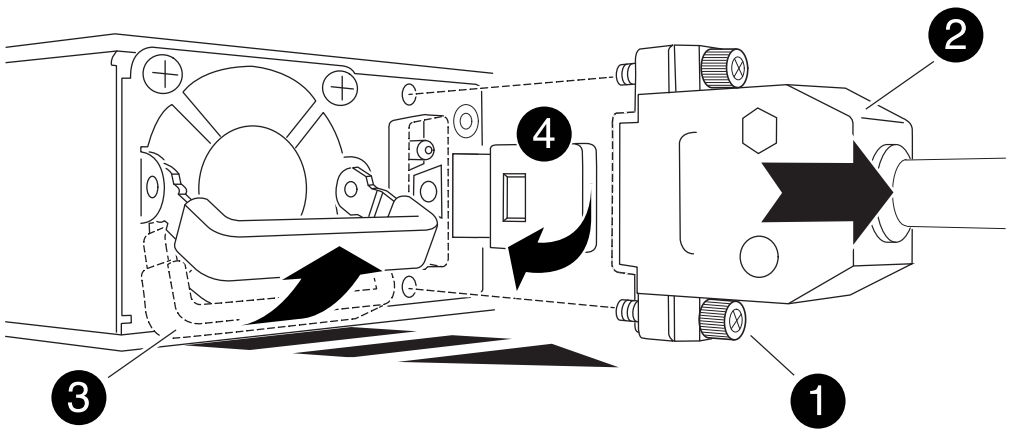

1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

5. Insert the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently slide the PSU into the controller until the locking tab clicks into place.

      A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.
6. Reconnect the D-SUB DC power cord:

   Once power is restored to the PSU, the status LED should be green.

   a. Plug the D-SUB DC power cord connector into the PSU.
   b. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A20, AFF A30, and AFF A50

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your AFF A20, AFF A30, or AFF A50 storage system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

About this task

- You can use this procedure with all versions of ONTAP supported by your storage system.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

   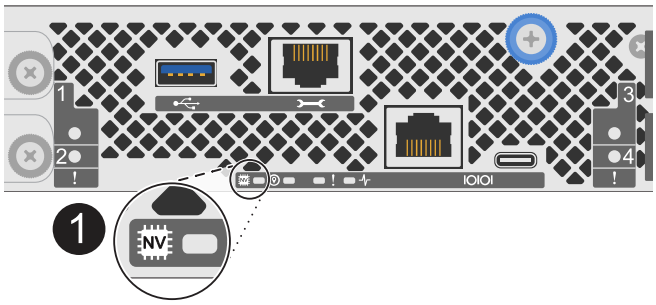

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

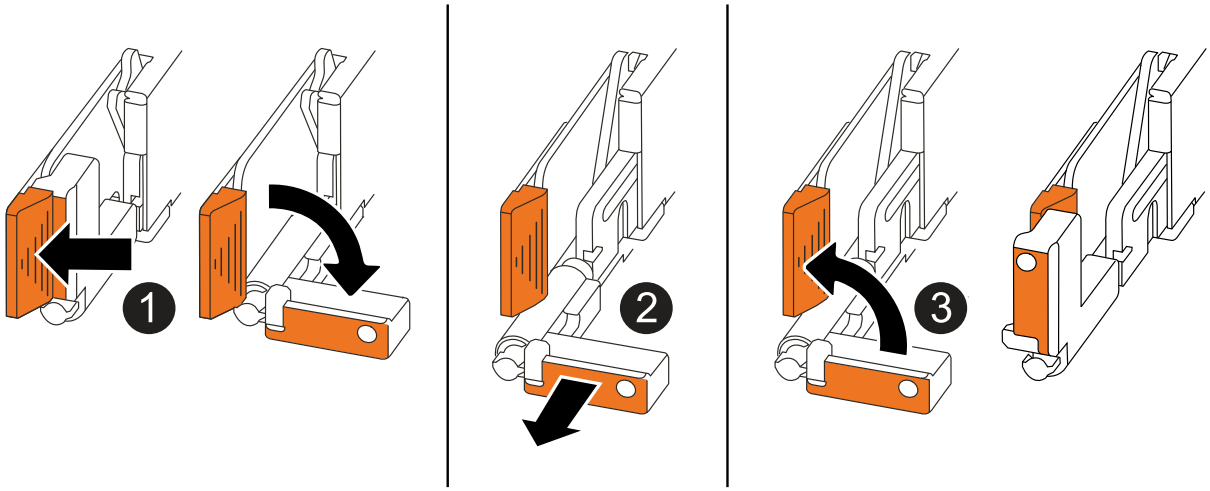

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the RTC battery

Remove the failed RTC battery and install the replacement RTC battery.

Steps

1. Locate the RTC battery.
2. Remove the RTC battery:

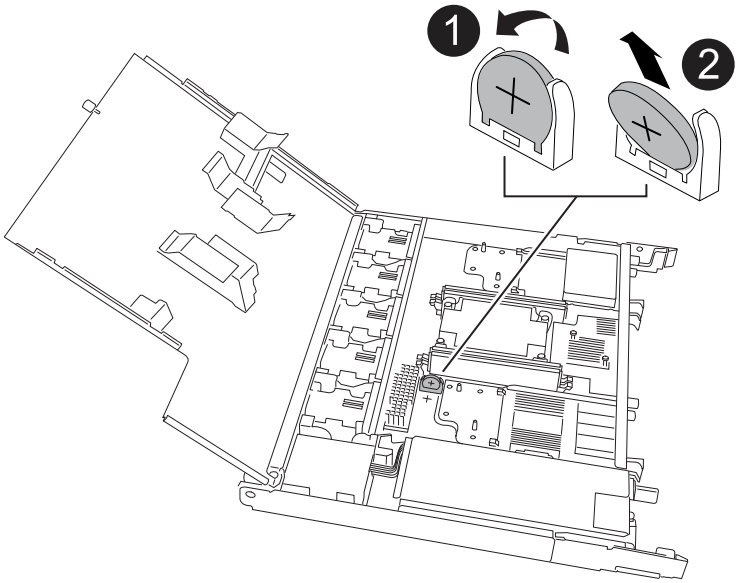

1	Gently rotate the RTC battery at an angle away from its holder.
2	Lift the RTC battery out of its holder.

3. Install the replacement RTC battery:
   a. Remove the replacement battery from the antistatic shipping bag.
   b. Position the battery so that the plus sign on the battery faces out to correspond with the plus sign on the motherboard.
   c. Insert the battery into the holder at an angle, and then push it into an upright position so it is fully seated in the holder.
   d. Visually inspect the battery to make sure that it is completely seated in its holder and that the polarity is correct.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

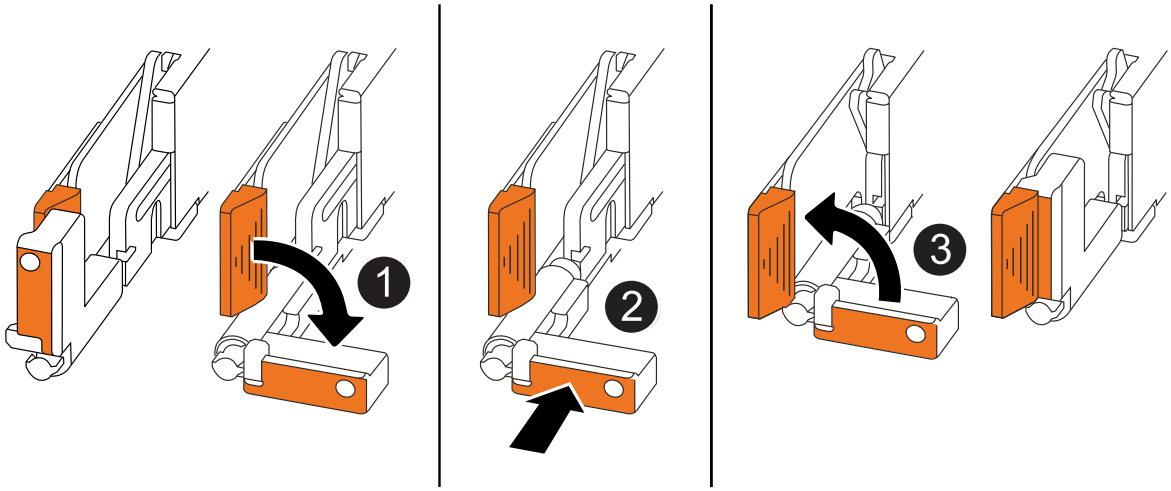

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Reset the time and date on the controller

> After replacing the RTC battery, inserting the controller and powering on first BIOS reset, you will see the following error messages:
> `RTC date/time error. Reset date/time to default`
> `RTC power failure error`
> These messages are expected and you can continue with this procedure.

1. On the healthy controller, check the date and time:

   `cluster date show`

> If your storage system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*.

2. On the impaired controller, at the LOADER prompt, check the time and date:

   `cluster date show`

   a. If necessary, modify the date:

      `set date mm/dd/yyyy`

   b. If necessary, set the time, in GMT:

```
set time hh:mm:ss
```

c. Confirm the date and time.

3. At the LOADER prompt, enter `bye` to reinitialize the I/O modules, other components, and let the controller reboot.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF A150 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

Quick guide - AFF A150

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

The Installation and Setup instructions give graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Use the xref:./a150/AFF A150 System Installation and Setup Instructions

The ASA A150 uses the same installation procedure as the AFF A150 system.

Video steps - AFF A150

The following video shows how to install and cable your system.

Animation - Install and setup of an AFF A150

If you have a MetroCluster configuration, use the MetroCluster documentation.

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

Detailed guide - AFF A150

This section gives detailed step-by-step instructions for installing an AFF A150 system.

If you have a MetroCluster configuration, use the MetroCluster documentation.

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

Step 1: Prepare for installation

To install your AFF A150 system, you create an account on the NetApp Support Site, register your system, and obtain your license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

- Make sure you have access to NetApp Hardware Universe (HWU) for information about site requirements as well as additional information on your configured system.
- Make sure you have access to the Release Notes for your version of ONTAP for more information about this system.
- Contact your network administrator for information about connecting your system to the switches.
- Make sure you have the following items at your site:
  - Rack space for the storage system
  - Phillips #2 screwdriver
  - Additional networking cables to connect your system to your network switch and laptop or console with a Web browser
  - A laptop or console with an RJ-45 connection and access to a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register your system.
4. Download and install Config Advisor on your laptop.
5. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
10 GbE cable (order dependent)	X6566B-05-R6 (112-00297), 0.5m X6566B-2-R6 (112-00299), 2m		Cluster interconnect network
10 GbE cable (order dependent)	Part number X6566B-2-R6 (112-00299), 2m or X6566B-3-R6 (112-00300), 3m X6566B-5-R6 (112-00301), 5m		Data
Optical network cables (order dependent)	X6553-R6 (112-00188), 2m X6536-R6 (112-00090), 5m X6554-R6(112-00189), 15m		FC host network
Cat 6, RJ-45 (order dependent)	Part numbers X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network and Ethernet data
Storage (order dependent)	Part number X66030A (112-00435), 0.5m X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		Storage
Micro-USB console cable	Not applicable		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	Not applicable		Powering up the system

6. Download and complete the Cluster Configuration Worksheet.

Step 2: Install the hardware

You install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

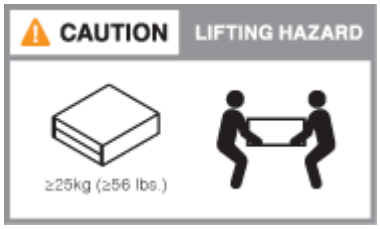

3. Attach cable management devices (as shown).

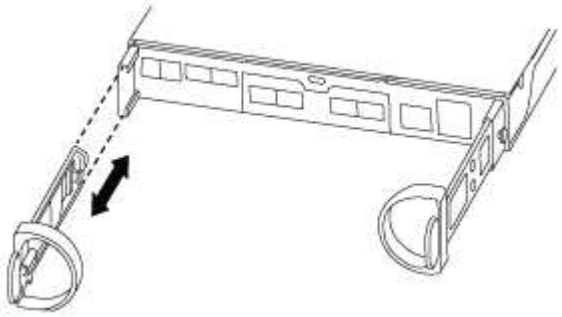

4. Place the bezel on the front of the system.

Step 3: Cable controllers to network

You cable the controllers to your network by using either the two-node switchless cluster method or the switched cluster method.

About this task

The following table identifies the cable type with the call out number and cable color in the illustrations for both two-node switchless cluster network cabling and switched cluster network cabling.

Cabling	Connection type
1	Cluster interconnect
2	Controllers to host data network switches
3	Controllers to management network switch

Option 1: Two-node switchless cluster

Cable your two-node switchless cluster.

About this task

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

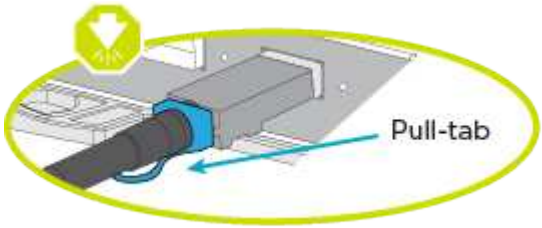

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Cable the cluster interconnect ports e0a to e0a and e0b to e0b with the cluster interconnect cable.

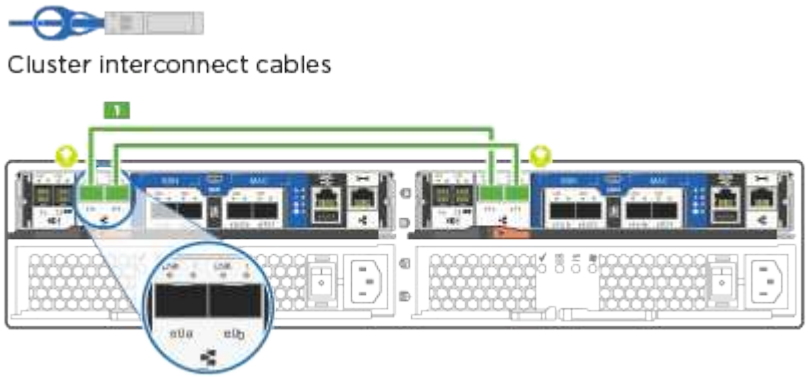

2. Cable the controllers to either a UTA2 data network or an Ethernet network:

UTA2 data network configurations

Use one of the following cable types to cable the UTA2 data ports to your host network.

- For an FC host, use 0c and 0d **or** 0e and 0f.
- For an 10GbE system, use e0c and e0d **or** e0e and e0f.

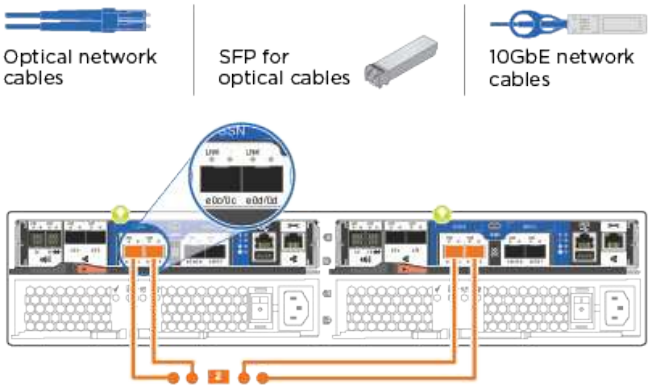

You can connect one port pair as CNA and one port pair as FC, or you can connect both port pairs as CNA or both port pairs as FC.

Ethernet network configurations

Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network.
in the following illustration.

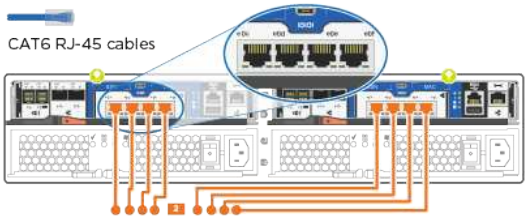

3. Cable the e0M ports to the management network switches with the RJ45 cables.

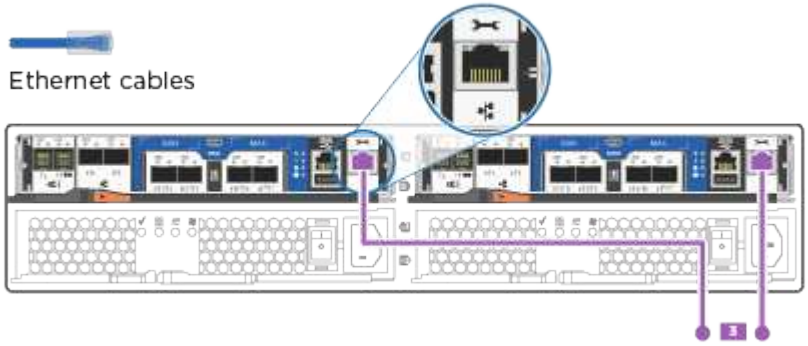

DO NOT plug in the power cords at this point.

Option 2: Switched cluster

Cable your switched cluster.

About this task

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

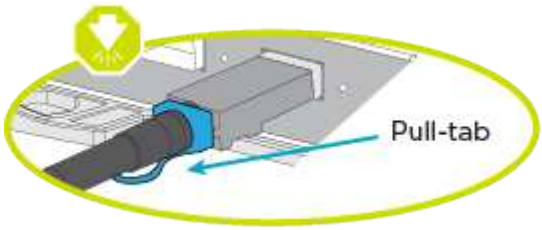

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. For each controller module, cable e0a and e0b to the cluster interconnect switches with the cluster interconnect cable.

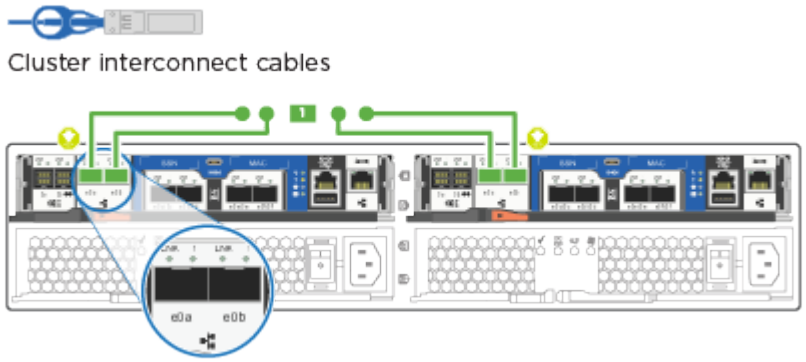

2. You can use either the UTA2 data network ports or the ethernet data network ports to connect the controllers to your host network:

UTA2 data network configurations

Use one of the following cable types to cable the UTA2 data ports to your host network.

- For an FC host, use 0c and 0d **or** 0e and 0f.
- For an 10GbE system, use e0c and e0d **or** e0e and e0f.

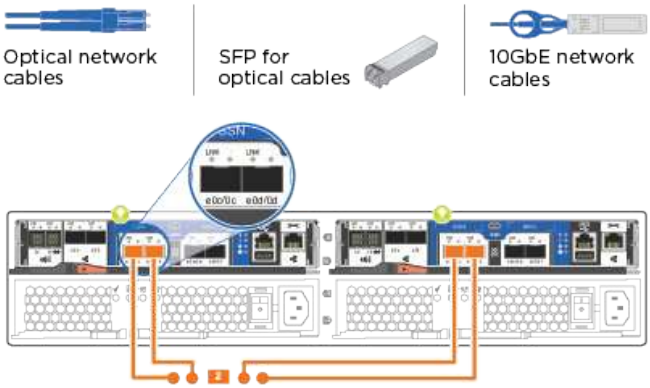

You can connect one port pair as CNA and one port pair as FC, or you can connect both port pairs as CNA or both port pairs as FC.

Ethernet network configurations

Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network.

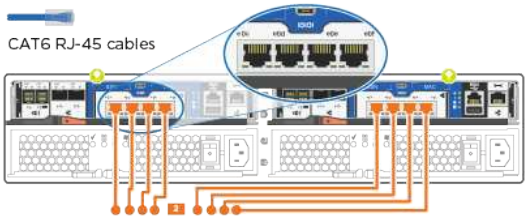

3. Cable the e0M ports to the management network switches with the RJ45 cables.

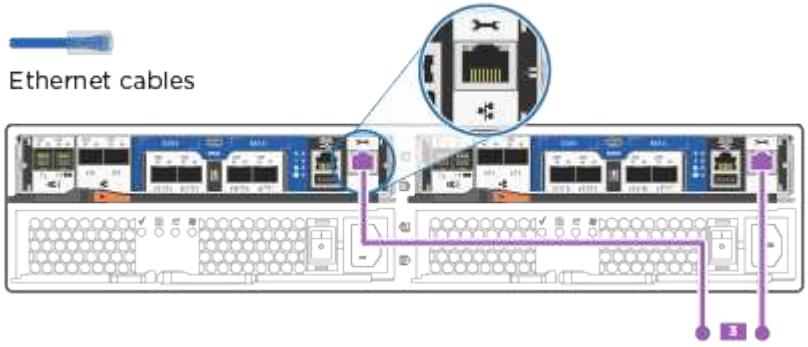

DO NOT plug in the power cords at this point.

Step 4: Cable controllers to drive shelves

Cable the controllers to your shelves using the onboard storage ports. NetApp recommends MP-HA cabling for systems with external storage.

About this task

- If you have a SAS tape drive, you can use single-path cabling. If you have no external shelves, MP-HA cabling to internal drives is optional (not shown) if the SAS cables are ordered with the system.
- You must cable the shelf-to-shelf connections, and then cable both controllers to the drive shelves.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

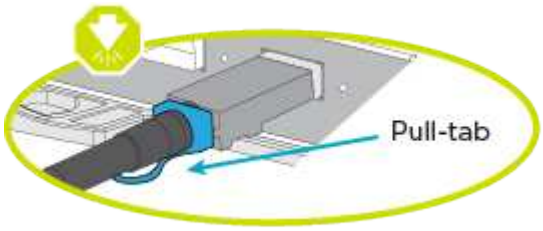

Steps

1. Cable the HA pair with external drive shelves.

   The following example shows cabling for DS224C drive shelves. The cabling is similar with other supported drive shelves.

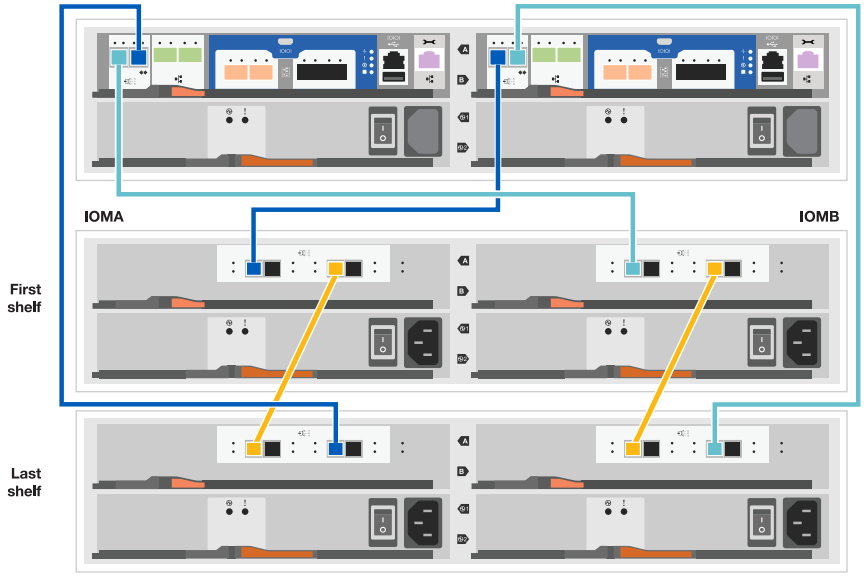

2. Cable the shelf-to-shelf ports.
   - Port 3 on IOM A to port 1 on the IOM A on the shelf directly below.
   - Port 3 on IOM B to port 1 on the IOM B on the shelf directly below.

 mini-SAS HD to mini-SAS HD cables

3. Connect each node to IOM A in the stack.
   - Controller 1 port 0b to IOM A port 3 on last drive shelf in the stack.
   - Controller 2 port 0a to IOM A port 1 on the first drive shelf in the stack.

 mini-SAS HD to mini-SAS HD cables

4. Connect each node to IOM B in the stack
   - Controller 1 port 0a to IOM B port 1 on first drive shelf in the stack.
   - Controller 2 port 0b to IOM B port 3 on the last drive shelf in the stack.

 mini-SAS HD to mini-SAS HD cables

For additional cabling information, see Install and cable shelves for a new system installation - shelves with IOM12/IOM12B modules.

Step 5: Complete system setup

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: If network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Use the following animation to set one or more drive shelf IDs

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Turn on the power switches to both nodes.

Initial booting may take up to eight minutes.

4. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

5. Connect your laptop to the Management switch.

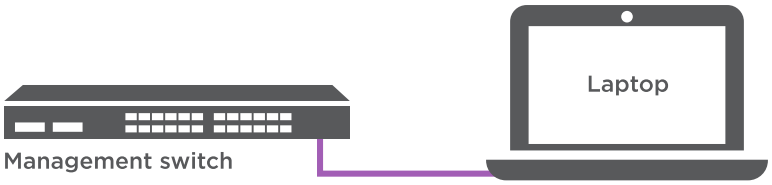

6. Select an ONTAP icon listed to discover:

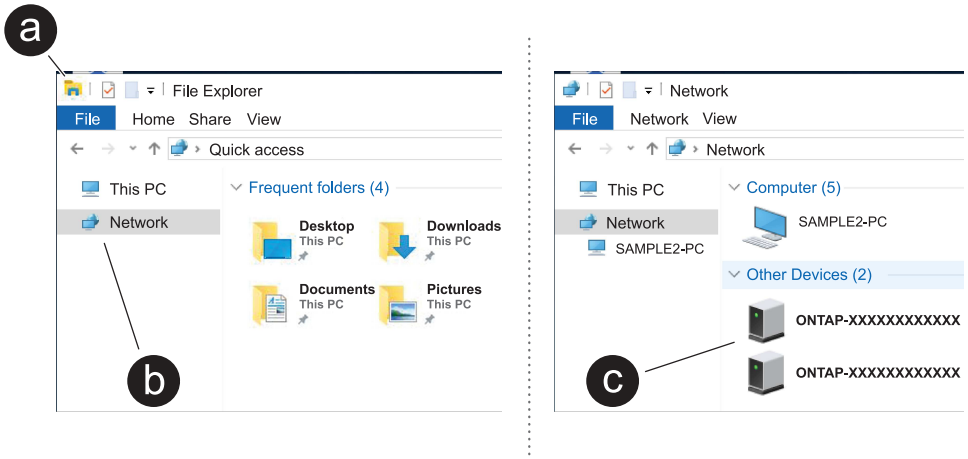

a. Open File Explorer.
b. Click **Network** in the left pane and right-click and select **refresh**.
c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

XXXXX is the system serial number for the target node.

System Manager opens.

7. Configure the system using the data you collected in the ONTAP Configuration Guide.
8. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create and account.
   b. Register your system.
   c. Download Active IQ Config Advisor.
9. Verify the health of your system by running Config Advisor.
10. After you have completed the initial configuration, go to the ONTAP documentation site for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console.
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for instructions on how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your system.

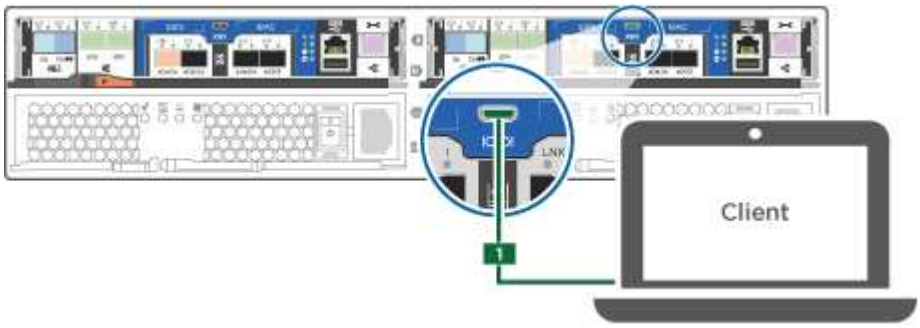

   c. Connect the laptop or console to the switch on the management subnet.

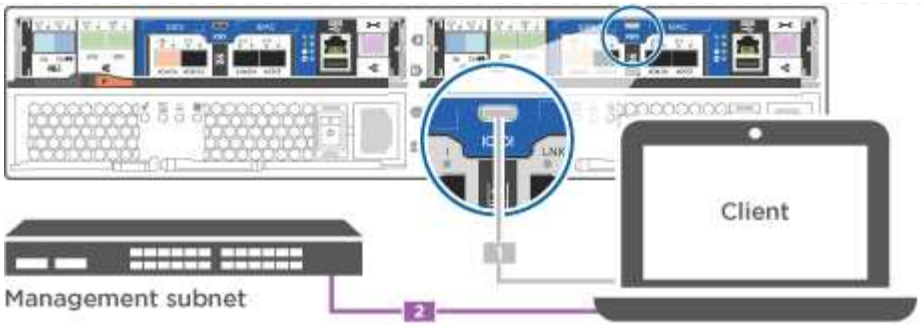

d. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Use the following animation to set one or more drive shelf IDs:

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
4. Turn on the power switches to both nodes.

Initial booting may take up to eight minutes.

5. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

6. Using System Manager on your laptop or console, configure your cluster.
   a. Point your browser to the node management IP address.

The format for the address is https://x.x.x.x.

b. Configure the system using the data you collected in the ONTAP Configuration Guide.

7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create and account.
   b. Register your system.
   c. Download Active IQ Config Advisor.
8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to the ONTAP documentation site for information about configuring additional features in ONTAP.

Maintain

Maintain AFF A150 hardware

For the AFF A150 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

NVEM Battery

A battery is included with a controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF A150

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct node:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy node* is the HA partner of the impaired node.

Check encryption key support and status - AFF A150

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:

- If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
- If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the impaired controller - AFF A150

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller displays…	Then…
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - AFF A150

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

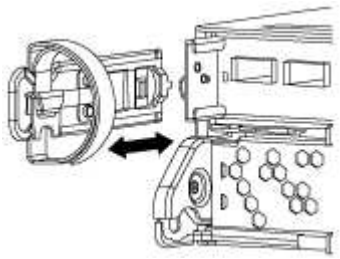

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

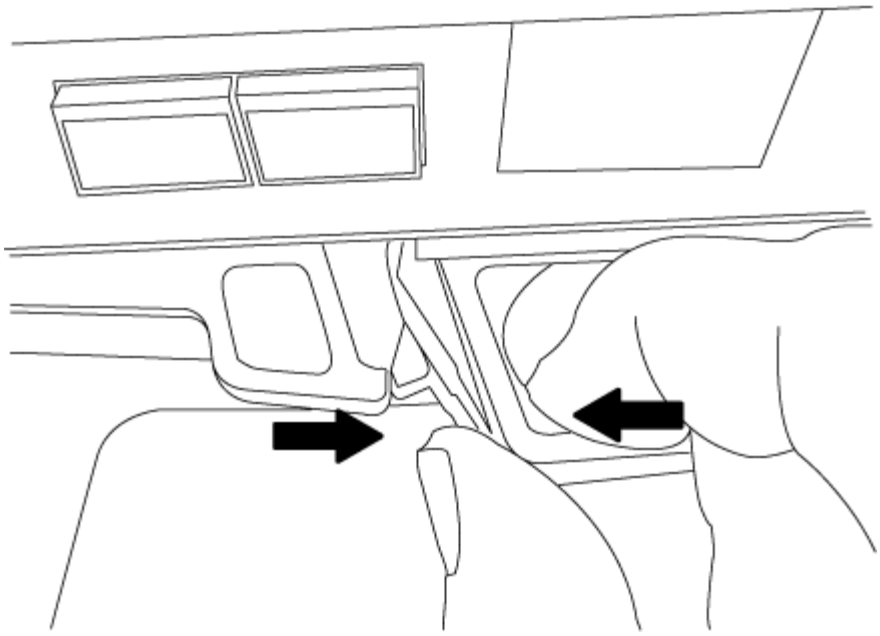

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

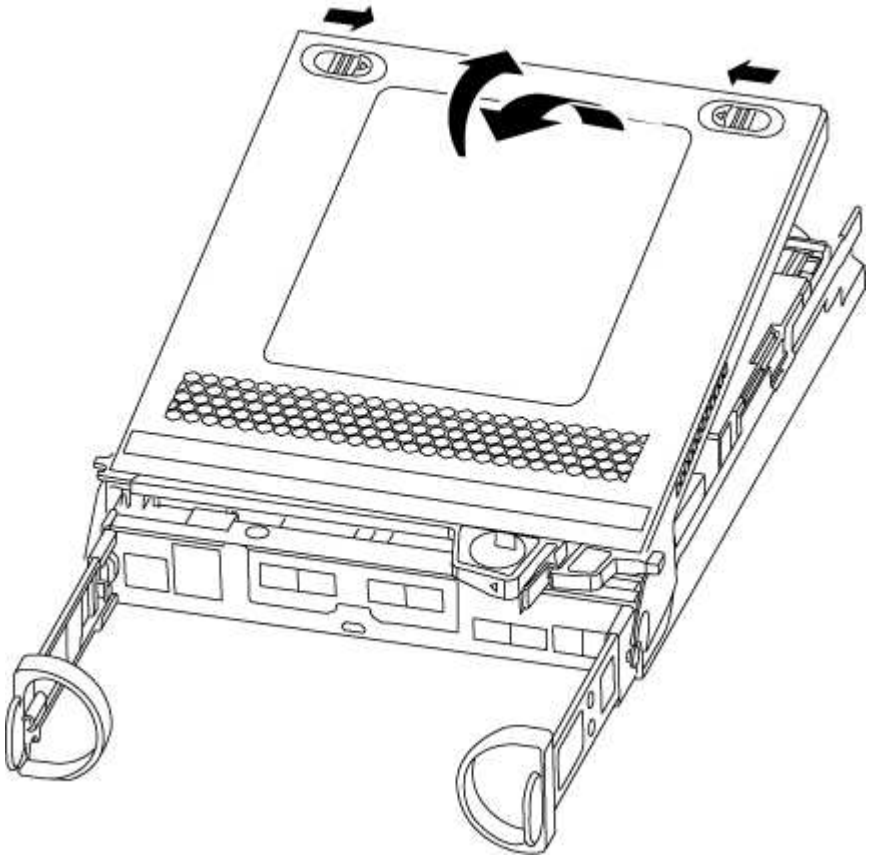

Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the boot media using the following illustration or the FRU map on the controller module:

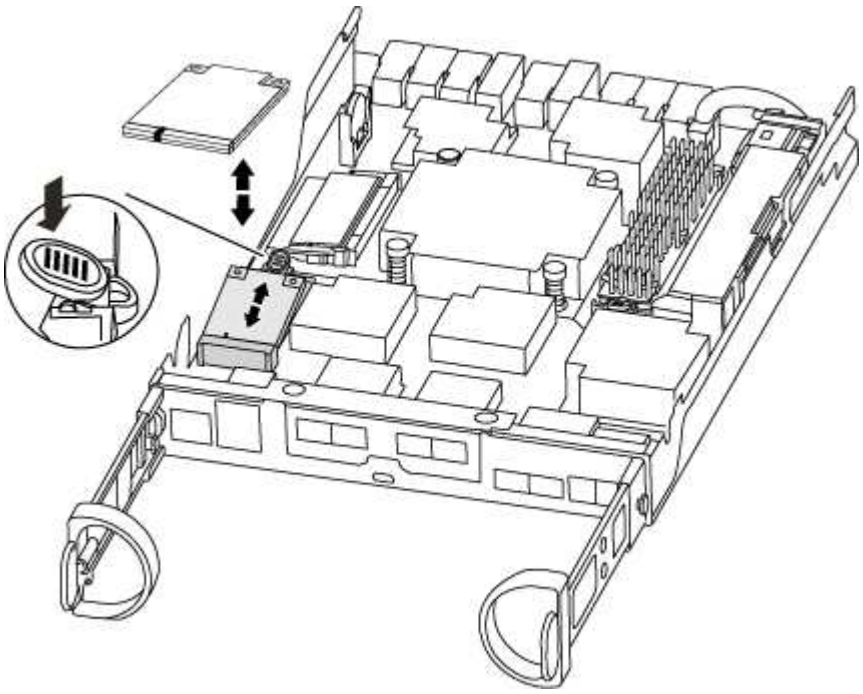

3. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

4. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
5. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

6. Push the boot media down to engage the locking button on the boot media housing.
7. Close the controller module cover.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

   The controller begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

6. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

   The system begins to boot and stops at the LOADER prompt.

7. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

     The target port you configure is the target port you use to communicate with the impaired controller from the healthy controller during var file system restore with a network connection. You can also use the e0M port in this command.

   - If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
     - filer_addr is the IP address of the storage system.
     - netmask is the network mask of the management network that is connected to the HA partner.
     - gateway is the gateway for the network.
     - dns_addr is the IP address of a name server on your network.
     - dns_domain is the Domain Name System (DNS) domain name.

       If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

Boot the recovery image - AFF A150

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:

- If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed
- If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.
8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore encryption - AFF A150

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.

- Backup information for the Onboard Key Manager.

- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF A150

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF A150

To replace the chassis, you must move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving all drives and controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-controller cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A150

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Controller is in a MetroCluster configuration

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Move and replace hardware - AFF A150

Move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Move a power supply

Moving out a power supply when replacing a chassis involves turning off, disconnecting, and removing the power supply from the old chassis and installing and connecting it on the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.
4. Use the cam handle to slide the power supply out of the system.

   

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

Step 2: Remove the controller module

Remove the controller module or modules from the old chassis.

1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

2. Remove and set aside the cable management devices from the left and right sides of the controller module.

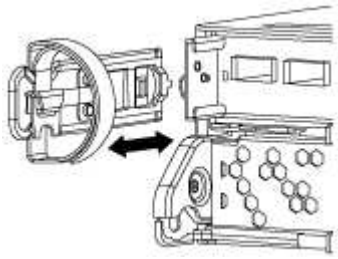

3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

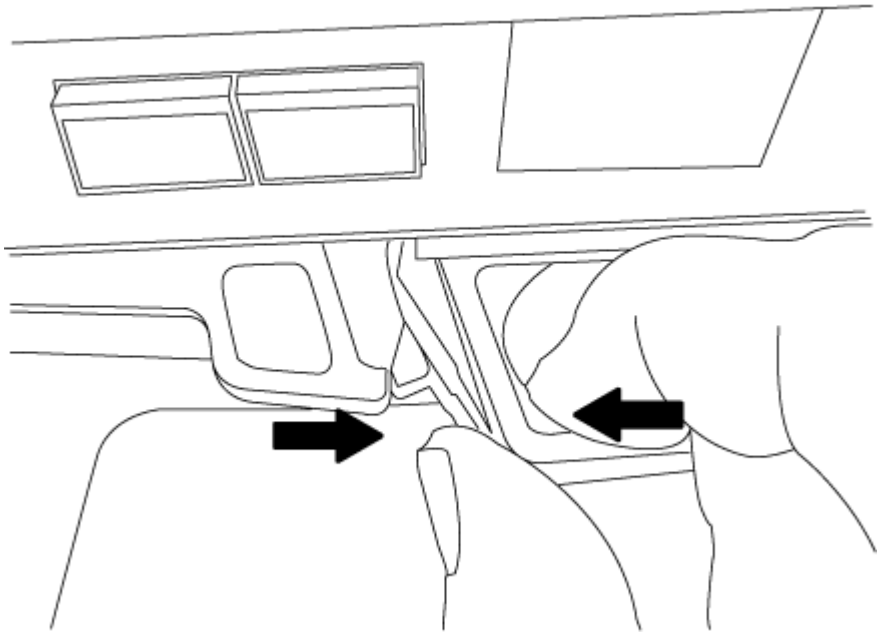

4. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 3: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new

chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.
5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It click when it is secure.
6. Repeat the process for the remaining drives in the system.

Step 4: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller

After you install the controller module and any other components into the new chassis, boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Repeat the preceding steps if there is a second controller to install in the new chassis.
4. Complete the installation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Repeat the preceding steps for the second controller module in the new chassis.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reinstall the blanking panel and then go to the next step.

5. Connect the power supplies to different power sources, and then turn them on.
6. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

> If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - AFF A150

You must verify the HA state of the chassis, switch back aggregates, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mcc-2n`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. Reboot the system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - AFF A150

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- This procedure includes steps for automatically or manually reassigning drives to the *replacement* controller, depending on your system’s configuration.

  You should perform the drive reassignment as directed in the procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - AFF A150

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - AFF A150

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove controller module

To replace the controller module, you must first remove the old controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

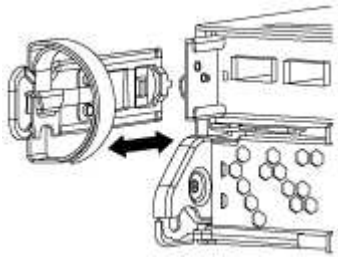

4. If you left the SFP modules in the system after removing the cables, move them to the new controller module.
5. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

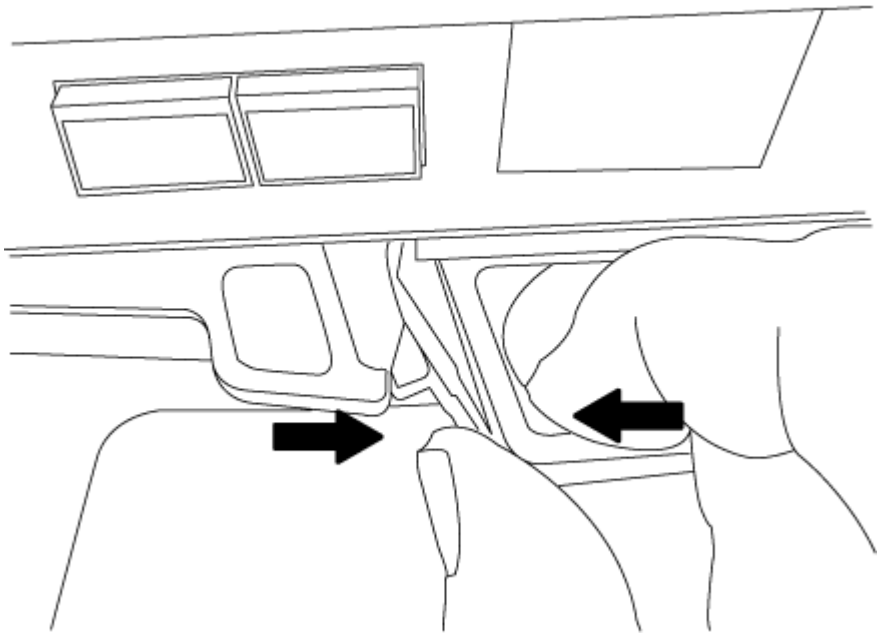

6. Turn the controller module over and place it on a flat, stable surface.
7. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

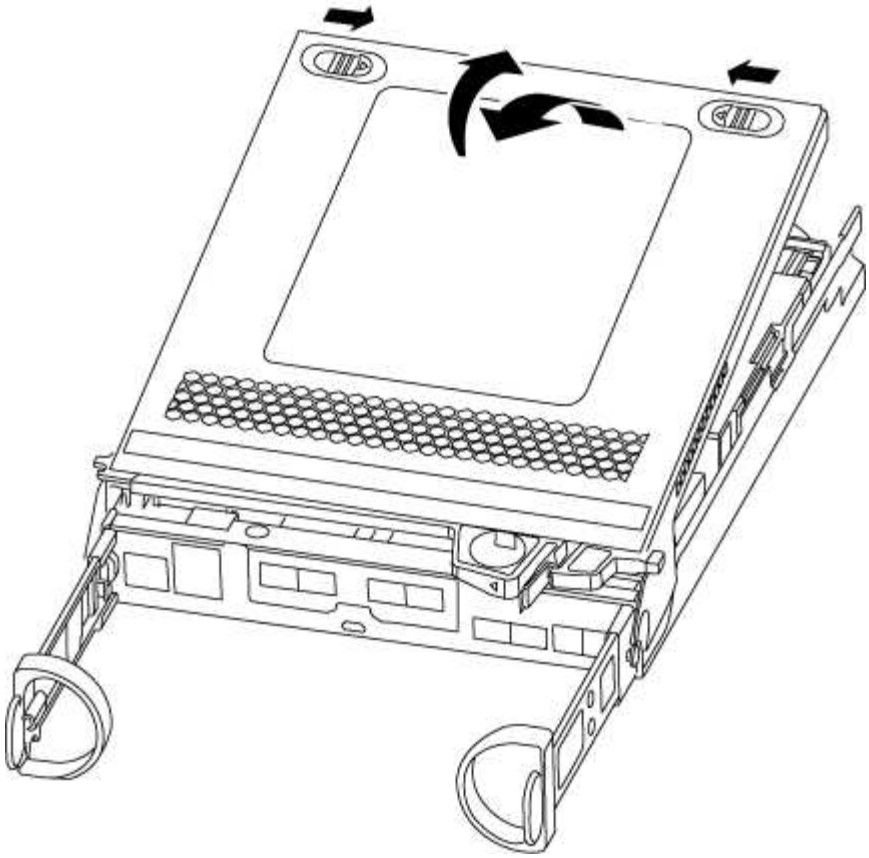

Step 2: Move the NVMEM battery

To move the NVMEM battery from the old controller module to the new controller module, you must perform a specific sequence of steps.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Locate the NVMEM battery in the controller module.

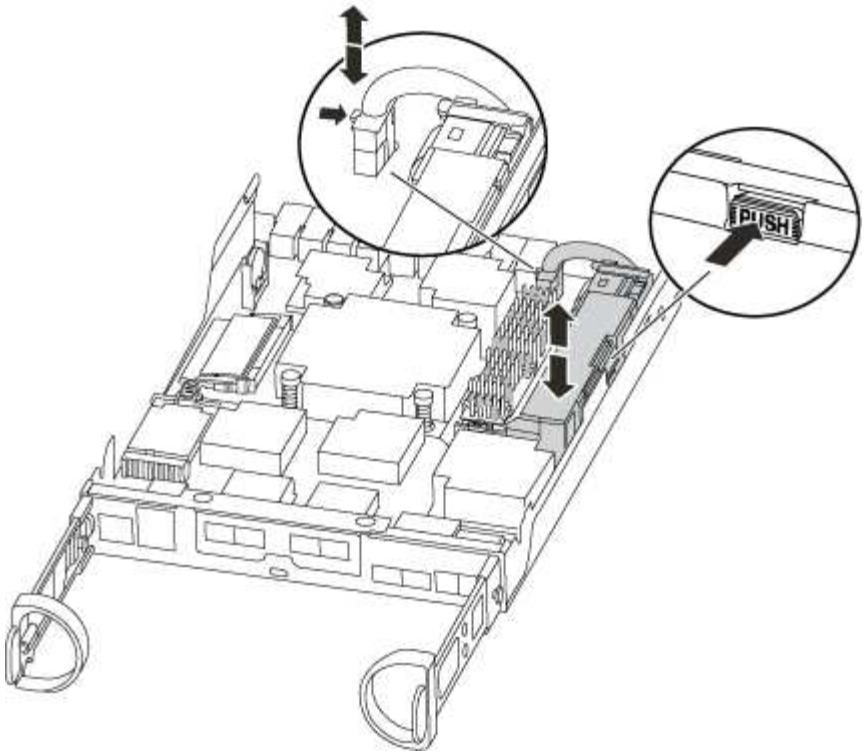

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Loop the battery cable around the cable channel on the side of the battery holder.
7. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
8. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.

Step 3: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller module and insert it in the new controller module.

1. Locate the boot media using the following illustration or the FRU map on the controller module:

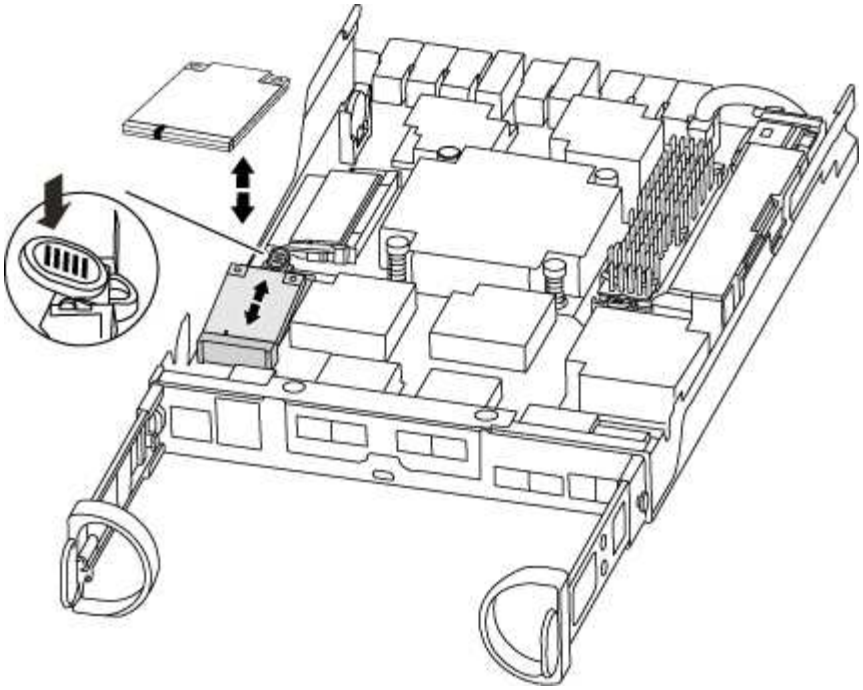

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 4: Move the DIMMs

To move the DIMMs, you must follow the directions to locate and move them from the old controller module into the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

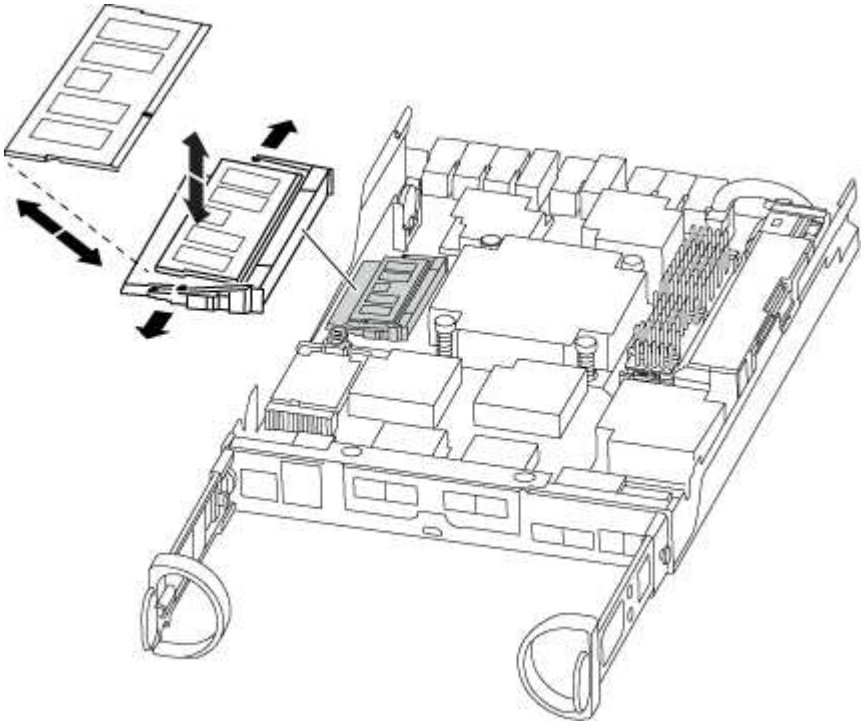

4. Repeat these steps to remove additional DIMMs as needed.
5. Verify that the NVMEM battery is not plugged into the new controller module.
6. Locate the slot where you are installing the DIMM.
7. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Repeat these steps for the remaining DIMMs.
9. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

   Make sure that the plug locks down onto the controller module.

Step 5: Move a caching module, if present

If your AFF A220 or FAS2700 system has a caching module, you need to move the caching module from the old controller module to the replacement controller module. The caching module is referred to as the “M.2 PCIe card” on the controller module label.

You must have the new controller module ready so that you can move the caching module directly from the old controller module to the corresponding slot in the new one. All other components in the storage system must be functioning properly; if not, you must contact technical support.

1. Locate the caching module at the rear of the controller module and remove it.

a. Press the release tab.

b. Remove the heatsink.

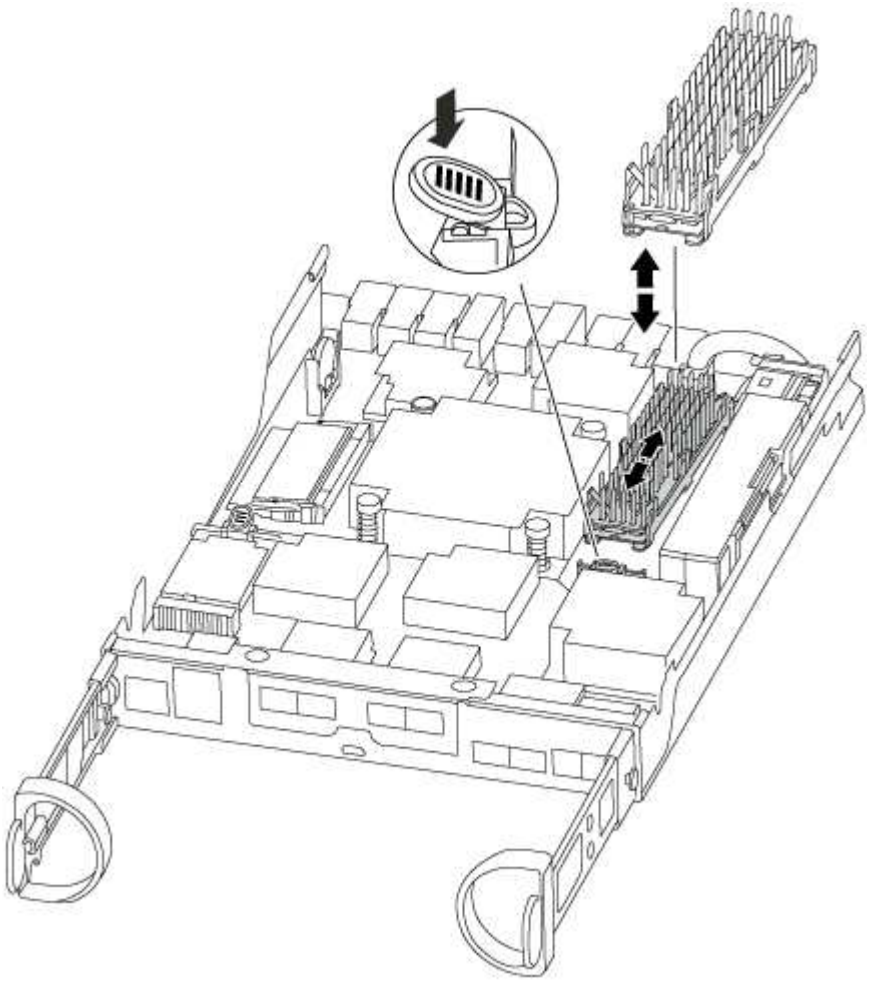

2. Gently pull the caching module straight out of the housing.
3. Move the caching module to the new controller module, and then align the edges of the caching module with the socket housing and gently push it into the socket.
4. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

5. Reseat and push the heatsink down to engage the locking button on the caching module housing.
6. Close the controller module cover, as needed.

Step 6: Install the controller

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Interrupt the boot process **only** after determining the correct timing: You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. Only press `Ctrl-C` when you see the message `Press Ctrl-C for Boot Menu`. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. Select the option to boot to Maintenance mode from the displayed menu.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process. e. Interrupt the boot process **only** after determining the correct timing: You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. Only press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. f. From the boot menu, select the option for Maintenance mode.

Important: During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
  You can safely respond `y` to these prompts.

Restore and verify the system configuration - AFF A150

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

- `ha`
- `mcc`
- `mcc-2n`
- `mccip`
- `non-ha`

b. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF A150

To complete the replacement procedure and restore your system to full operation, you must recable the storage, confirm disk reassignment, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks.

You must use the correct procedure for your configuration:

Controller redundancy	Then use this procedure…
HA pair	Option 1: Verify the system ID change on an HA system
Stand-alone	Option 2: Manually reassign the system ID on a stand-alone system in ONTAP
Two-node MetroCluster configuration	Option 3: Manually reassign the system ID on systems in a two-node MetroCluster configuration

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *`local-node-name`* `partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

- Restore onboard key management encryption keys
- Restore external key management encryption keys

6. Give back the controller:

   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Option 2: Manually reassign the system ID on a stand-alone system in ONTAP

In a stand-alone system, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems that are in a stand-alone configuration.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by pressing Ctrl-C, and then select the option to boot to Maintenance mode from the displayed menu.
2. You must enter `Y` when prompted to override the system ID due to a system ID mismatch.
3. View the system IDs: `disk show -a`
4. You should make a note of the old system ID, which is displayed as part of the disk owner column.

   The following example shows the old system ID of 118073209:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                  POOL   SERIAL NUMBER  HOME
--------    -------------          -----  -------------  -------------
disk_name    system-1  (118073209)  Pool0  J8XJE9LC       system-1
(118073209)
disk_name    system-1  (118073209)  Pool0  J8Y478RC       system-1
(118073209)
.
.
.
```

5. Reassign disk ownership by using the system ID information obtained from the disk show command: `disk reassign -s old system ID disk reassign -s 118073209`
6. Verify that the disks were assigned correctly: `disk show -a`

   The disks belonging to the replacement node should show the new system ID. The following example now show the disks owned by system-1 the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                  POOL   SERIAL NUMBER  HOME
--------    -------------          -----  -------------  -------------
disk_name    system-1  (118065481)  Pool0  J8Y0TDZC       system-1
(118065481)
disk_name    system-1  (118065481)  Pool0  J8Y0TDZC       system-1
(118065481)
.
.
.
```

7. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

- Restore onboard key management encryption keys
- Restore external key management encryption keys

8. Boot the node: `boot_ontap`

Option 3: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `metrocluster node show -fields node-systemid,dr-partner-systemid`

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

   ```
   dr-group-id cluster         node                 node-systemid dr-
   partner-systemid
    ----------- --------------------- ------------------- -------------
   -------------------
    1           Cluster_A             Node_A_1             536872914
   118073209
    1           Cluster_B             Node_B_1             118073209
   536872914
    2 entries were displayed.
   ```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

   ```
   Local System ID: 118065481
       ...
       ...
   ```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                  POOL   SERIAL NUMBER  HOME
-------    -------------          -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Verify that the coredumps are saved: `system node run -node *local-node-name* partner savecore`

      If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node *local-node-name* partner savecore -s command`.</info>.

   c. Return to the admin privilege level: `set -privilege admin`

7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`

8. Boot the *replacement* node: `boot_ontap`

9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`

10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

       You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

    b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`
    c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - AFF A150

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.

4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A150

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

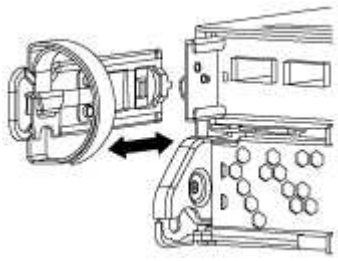

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

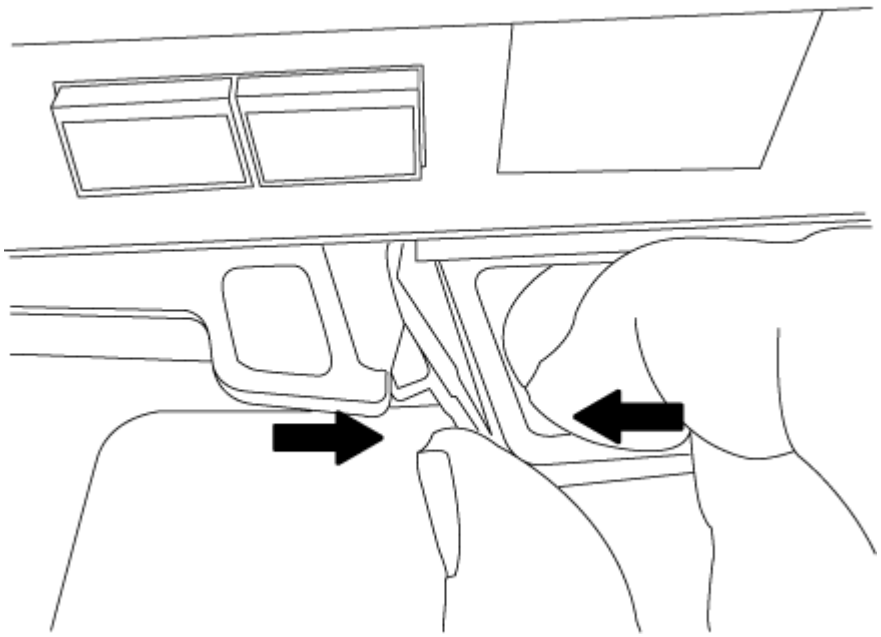

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

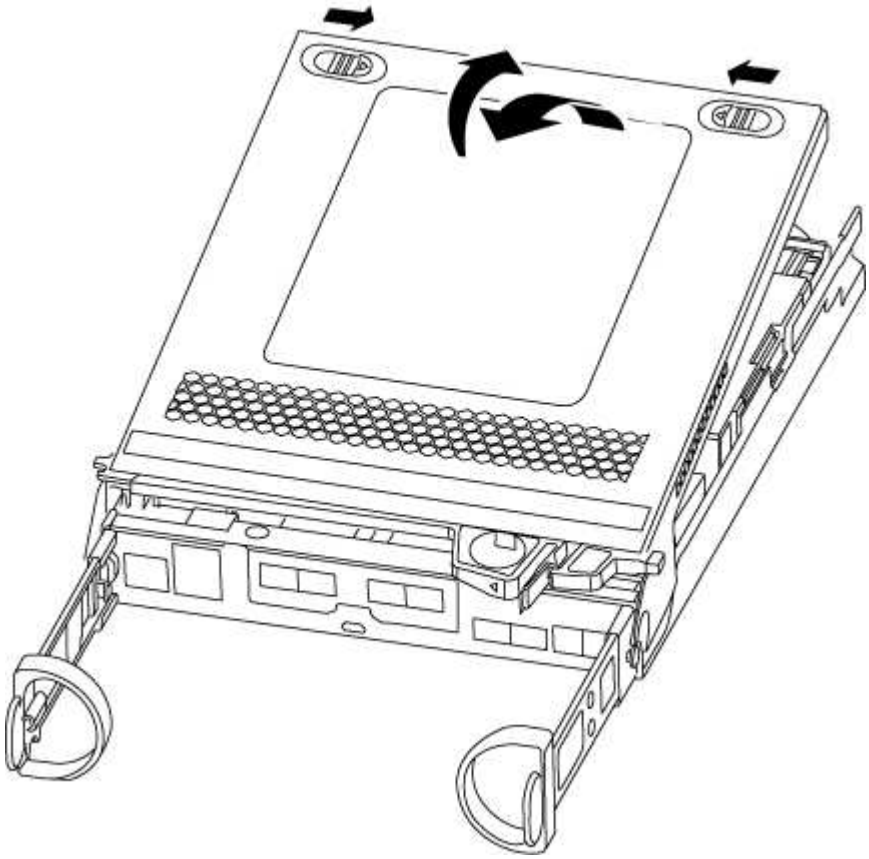

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

If you are replacing a DIMM, you need to remove it after you have unplugged the NVMEM battery from the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED on the back of controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

   

3. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
4. If the NVMEM LED is flashing, there is data in the NVMEM and you must disconnect the battery to clear the memory:
   a. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

   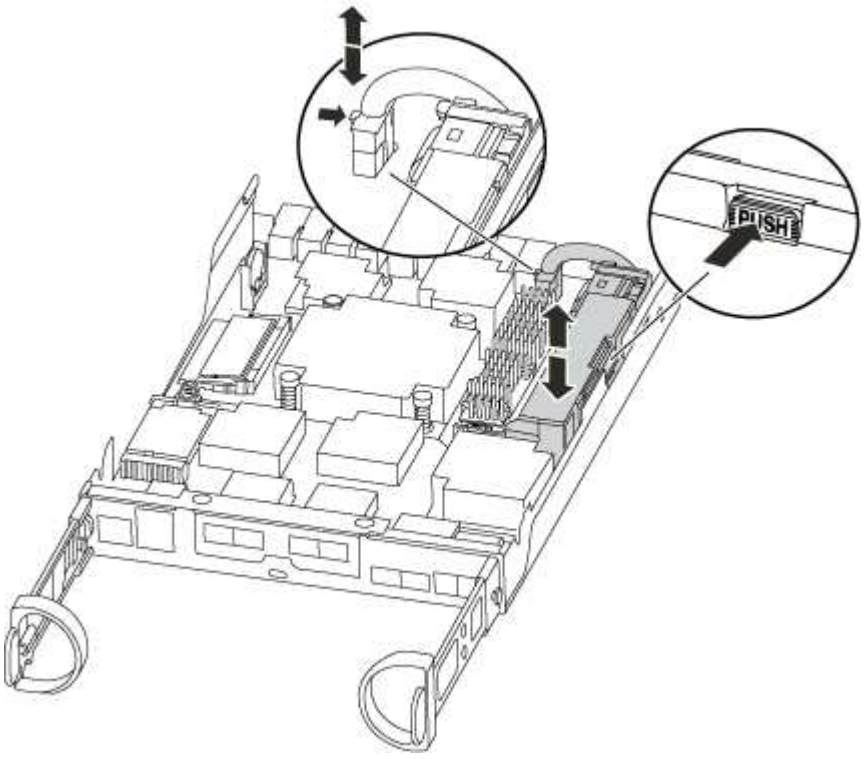

   b. Confirm that the NVMEM LED is no longer lit.
   c. Reconnect the battery connector.

5. Return to Step 3: Replace the DIMMs of this procedure to recheck the NVMEM LED.
6. Locate the DIMMs on your controller module.
7. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
8. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   The number and placement of system DIMMs depends on the model of your system.

   The following illustration shows the location of system DIMMs:

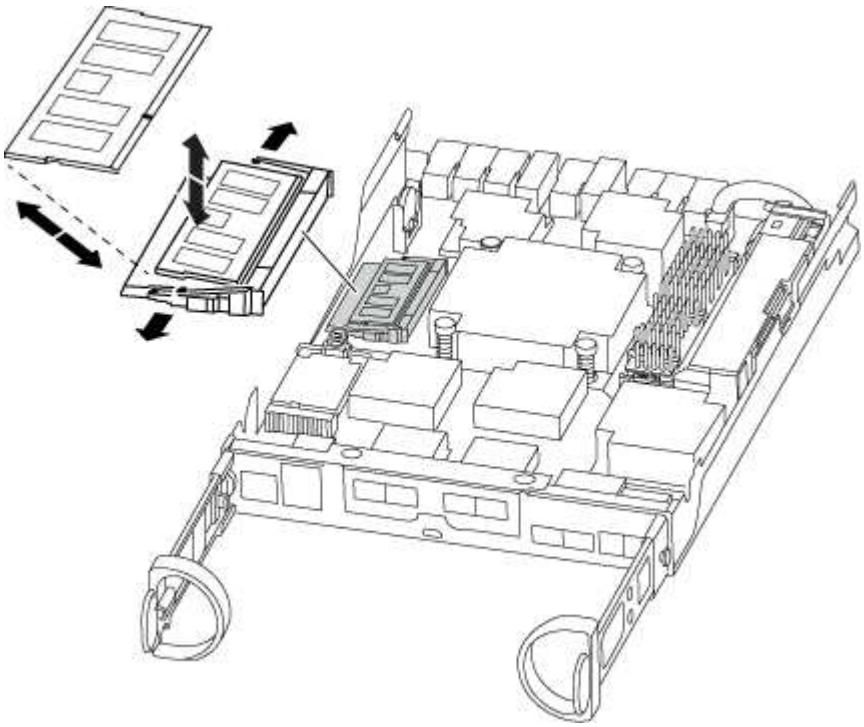

9. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

10. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

    The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

    Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

11. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
12. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

Make sure that the plug locks down onto the controller module.

13. Close the controller module cover.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, then turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - AFF A150

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace the NVMEM battery - AFF A150

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

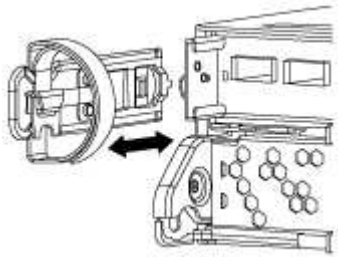

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

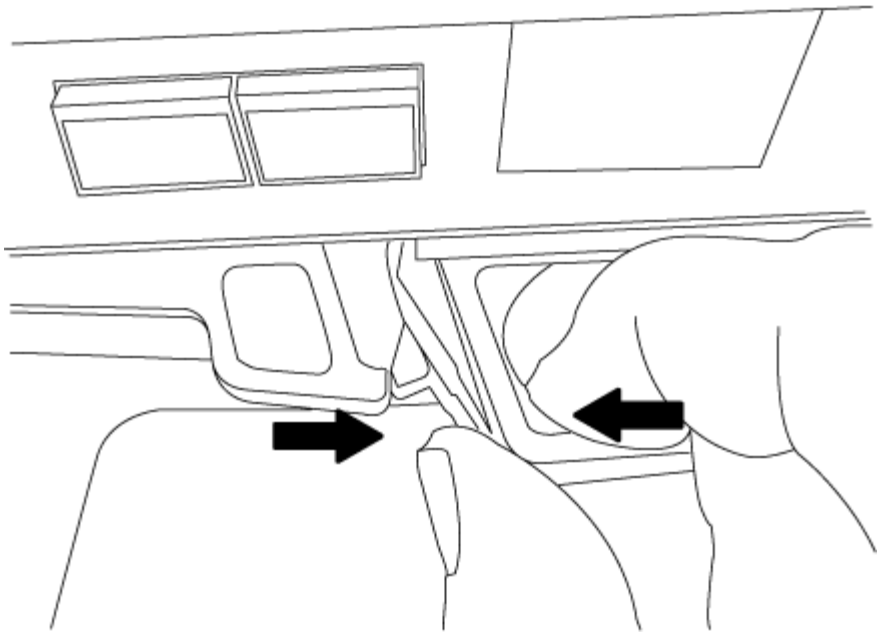

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery in your system, you must remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

3. Locate the NVMEM battery in the controller module.

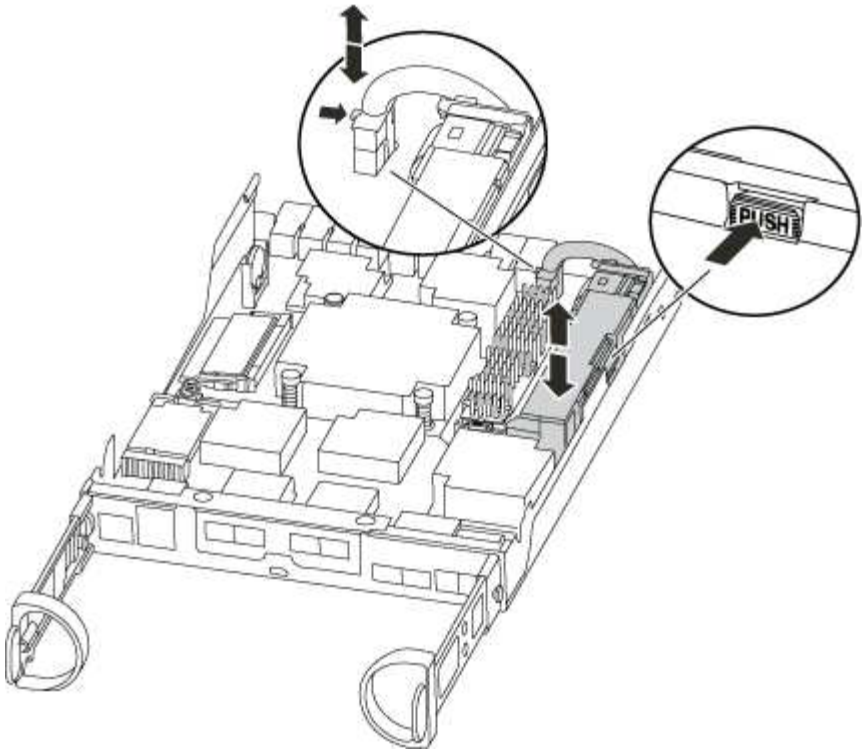

4. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
5. Remove the battery from the controller module and set it aside.
6. Remove the replacement battery from its package.
7. Loop the battery cable around the cable channel on the side of the battery holder.
8. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
9. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
10. Plug the battery plug back into the controller module.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber

optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - AFF A150

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> Cooling is integrated with the power supply, so you must replace the power supply within two minutes of removal to prevent overheating due to reduced airflow. Because the chassis provides a shared cooling configuration for the two HA nodes, a delay longer than two minutes will shut down all controller modules in the chassis. If both controller modules do shut down, make sure that both power supplies are inserted, turn both off for 30 seconds, and then turn both on.

- Power supplies are auto-ranging.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.

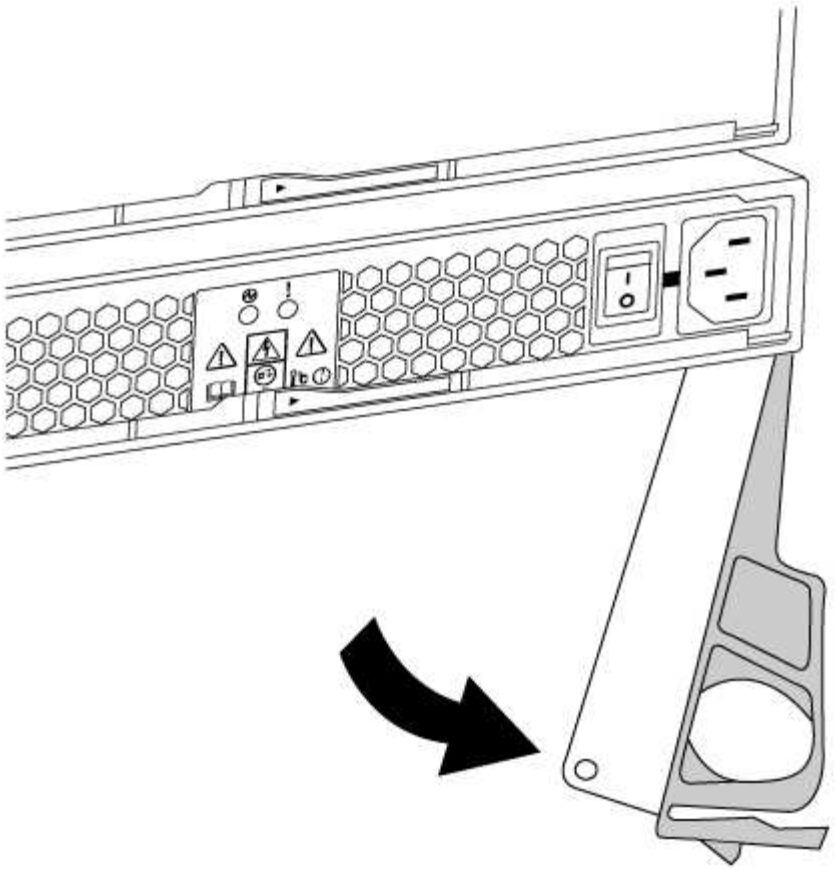

5. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.
7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
9. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.

10. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

    The power supply LEDs are lit when the power supply comes online.

11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A150

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

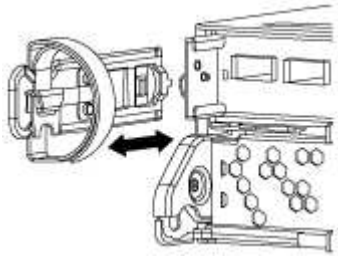

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

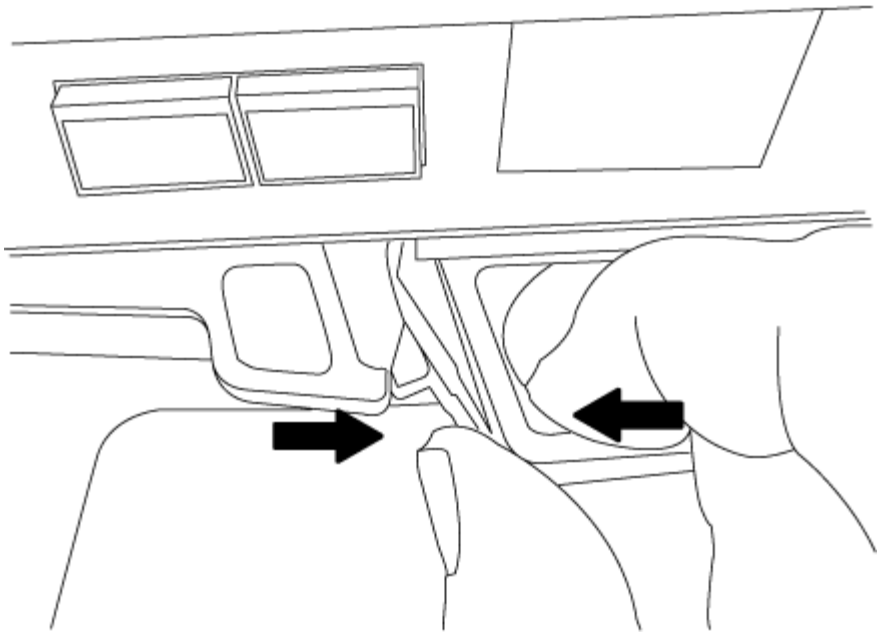

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

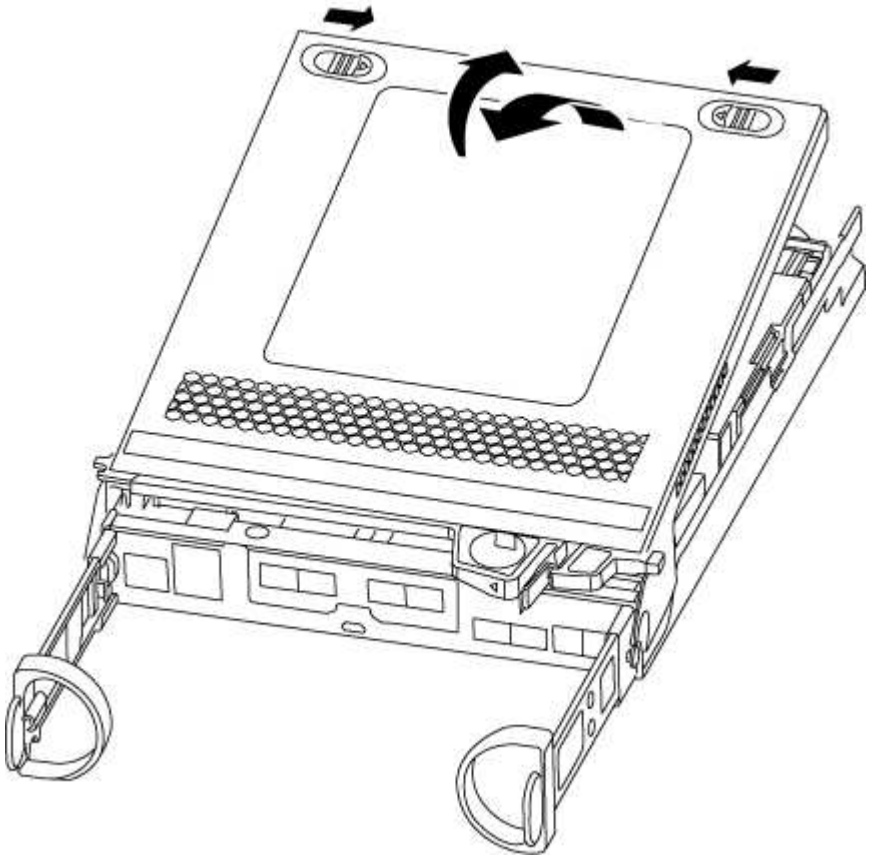

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

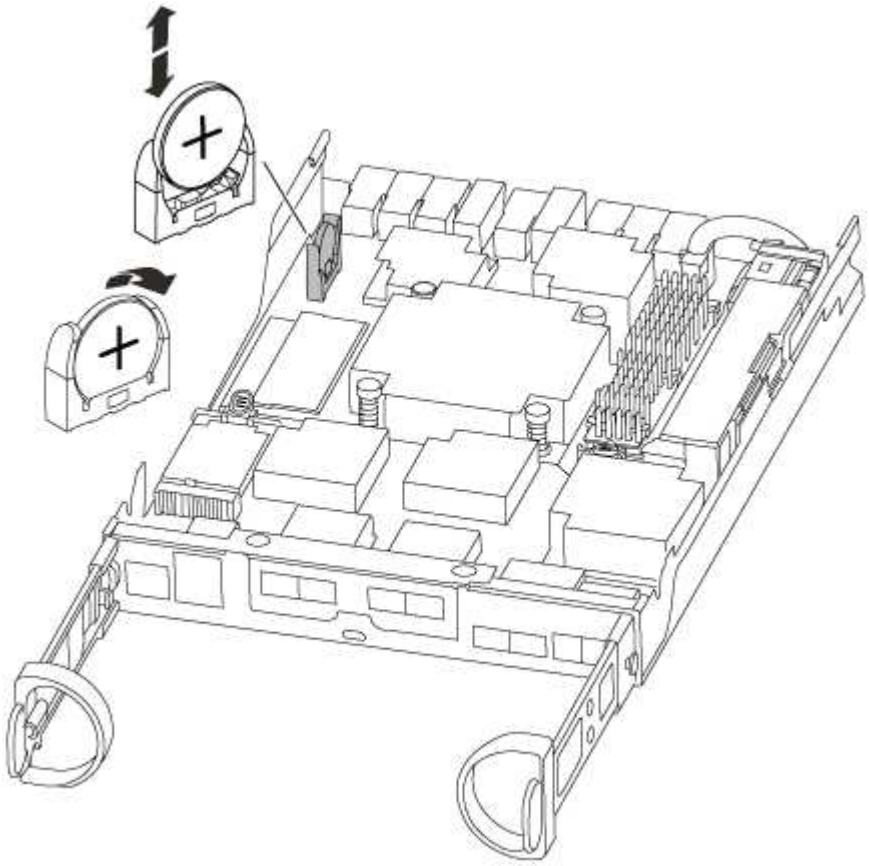

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and set time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF A250 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick steps - AFF A250

The Installation and Setup instructions give graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

The ASA A250 and ASA C250 use the same installation procedure as the AFF A250 system.

AFF A250 Installation and Setup Instructions

Video steps - AFF A250

The following video shows how to install and cable your new system.

Animation - Installation and Setup of an AFF A250

Detailed steps - AFF A250

This section gives detailed step-by-step instructions for installing an AFF A250 system.

Step 1: Prepare for installation

To install your AFF A250 system, you need to create an account and register the system. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Customers with specific power requirements must check HWU for their configuration options.

Before you begin

- Make sure you have access to the NetApp Hardware Universe (HWU) for information about site requirements.
- Make sure you have access to the Release Notes for your version of ONTAP for more information about this system.
- You need to provide the following at your site:
  - Rack space for the storage system
  - Phillips #2 screwdriver
  - Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register your system.
4. Download and install NetApp Downloads: Config Advisor on your laptop.
5. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For… .2+
25 GbE cable	X66240A-05 (112-00595), 0.5m; X66240-2 (112-00573), 2m .3+		Cluster interconnect network
X66240A-2 (112-00598), 2m; X66240A-5 (112-00600), 5m	Data	100 GbE cable	X66211-2 (112-00574), 2m; X66211-5 (112-00576), 5m
Storage	RJ-45 (order dependent)	Not applicable	

Type of cable…	Part number and length	Connector type	For… .2+
Management network (BMC and wrench port) and Ethernet data (e0a and e0b)	Fibre Channel	X66250-2 (112-00342) 2m; X66250-5 (112-00344) 5m; X66250-15 (112-00346) 15m; X66250-30 (112-00347) 30m	
	Micro-USB console cable	Not applicable	
Console connection during software setup	Power cables	Not applicable	

6. Review the ONTAP Configuration Guide and collect the required information listed in that guide.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

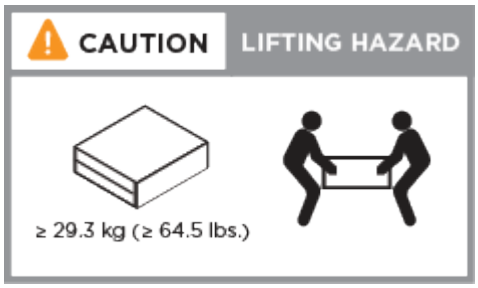

3. Identify and manage cables because this system does not have a cable management device.
4. Place the bezel on the front of the system.

Step 3: Cable controllers to cluster

Cable the controllers to a cluster by using the two-node switchless cluster method or by using the cluster

interconnect network method.

Option 1: Two-node switchless cluster

The management, Fibre Channel, and data or host network ports on the controller modules are connected to switches. The cluster interconnect ports are cabled on both controller modules.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

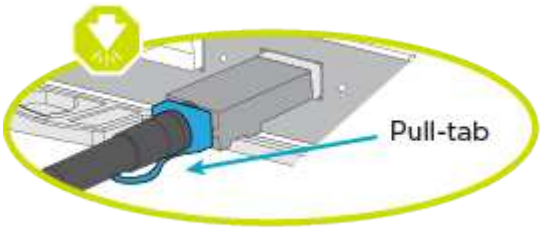

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the tabulated steps to complete the cabling between the controllers and the switches. Perform the steps on each controller.

Animation - Cable two-node switchless cluster

Steps

1. Use the the 25GbE cluster interconnect cable to connect the cluster interconnect ports e0c to e0c and e0d to e0d.


:

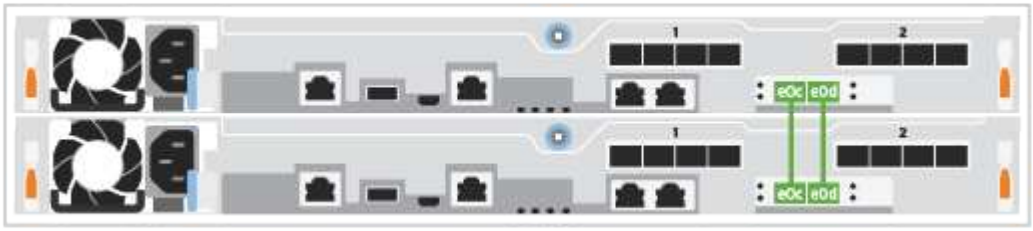

2. Cable the wrench ports to the management network switches with the RJ45 cables.

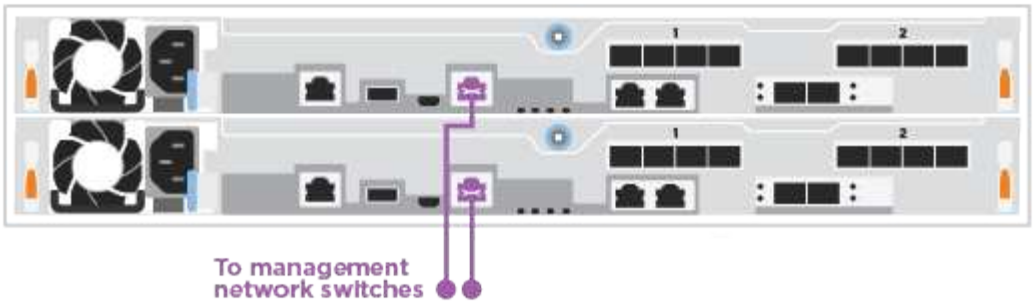

DO NOT plug in the power cords at this point.

Option 2: Switched cluster

All ports on the controllers are connected to switches; cluster interconnect, management, Fibre Channel, and data or host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

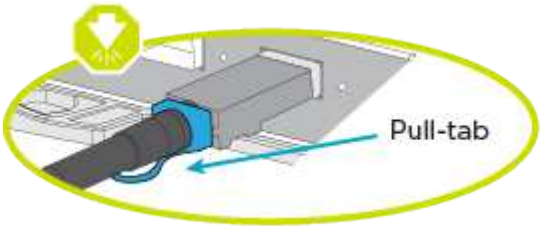

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the steps to complete the cabling between the controllers and the switches. Perform the steps on each controller.

Animation - Cable switched cluster

Steps

1. Cable the cluster interconnect ports e0c and e0d to the 25 GbE cluster interconnect switches.

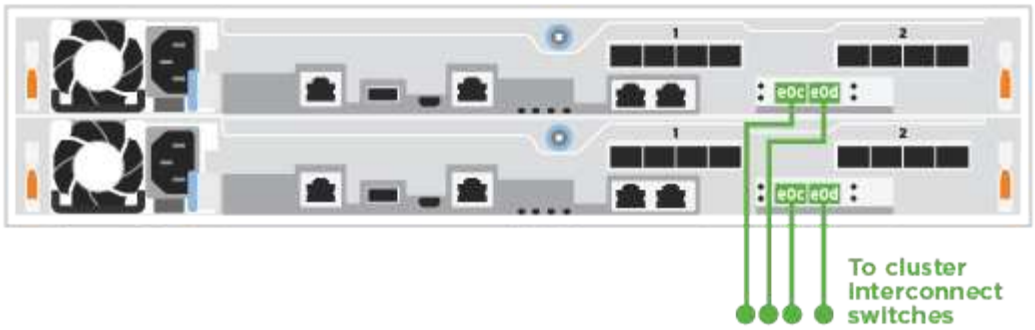

2. Cable the wrench ports to the management network switches with the RJ45 cables.

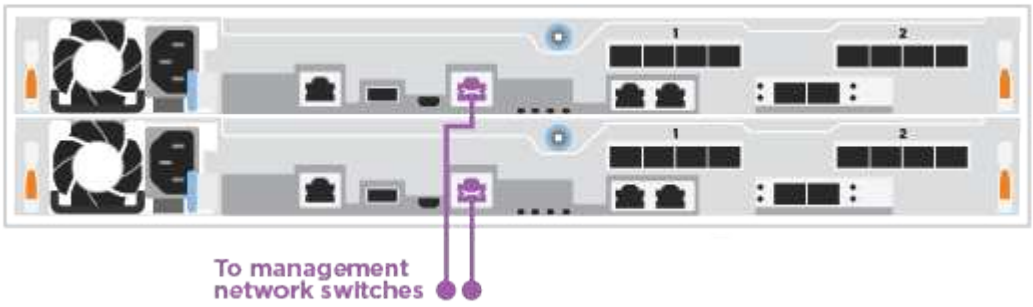

Step 4: Cable to host network or storage (Optional)

You have configuration-dependent optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cabling to a host network and storage.

NetApp Hardware Universe slot priority for host network cards (Fibre Channel or 25GbE) is slot 2. However, if you have both cards, the Fibre Channel card goes in slot 2 and the 25GbE card goes in slot 1 (as shown in the options below). If you have an external shelf, the storage card goes in slot 1, the only supported slot for shelves.

Option 1: Cable to Fibre Channel host network

Fibre Channel ports on the controllers are connected to Fibre Channel host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

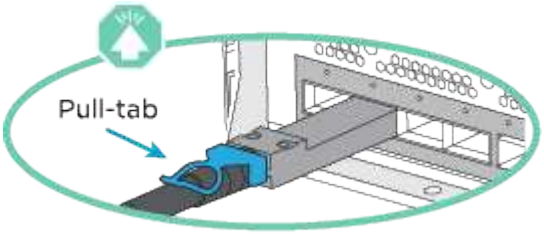

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again..

About this task

Perform the following step on each controller module.

Steps

1. Cable ports 2a through 2d to the FC host switches.

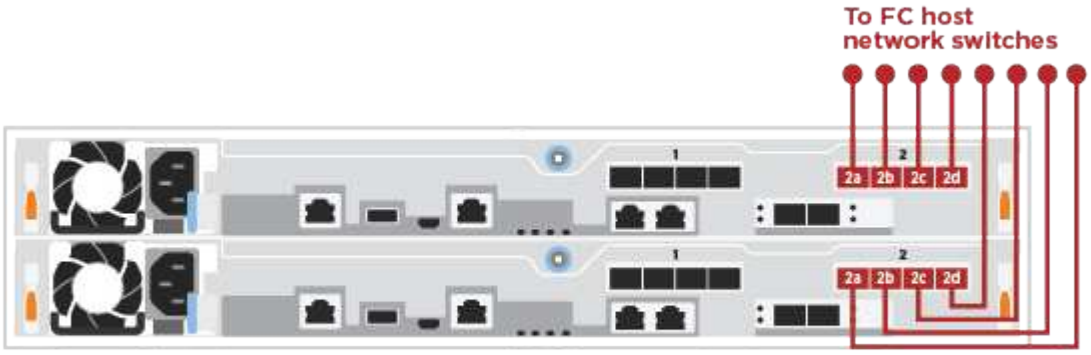

Option 2: Cable to 25GbE data or host network

25GbE ports on the controllers are connected to 25GbE data or host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

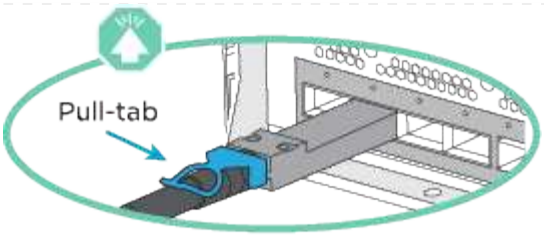

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Perform the following step on each controller module.

Steps

1. Cable ports e4a through e4d to the 10GbE host network switches.

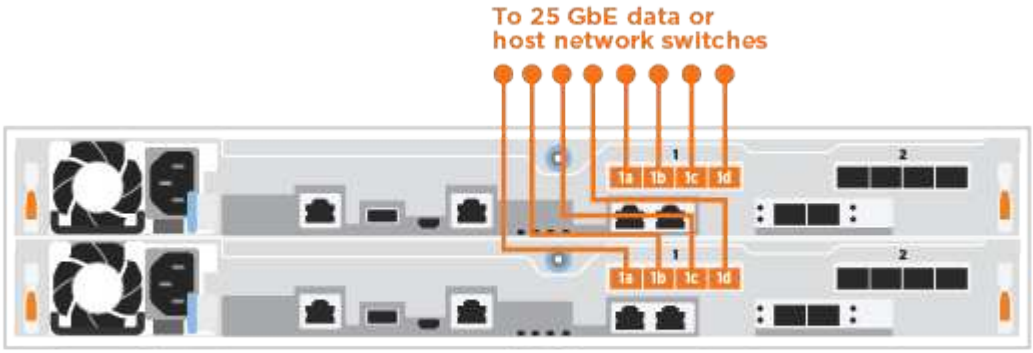

Option 3: Cable controllers to single drive shelf

Cable each controller to the NSM modules on the NS224 drive shelf.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

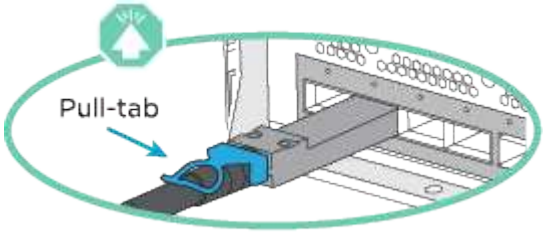

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the tabulated steps to complete the cabling between the controllers and the single shelf. Perform the steps on each controller module.

Animation - Cable the controllers to a single NS224

Steps

1. Cable controller A to the shelf.

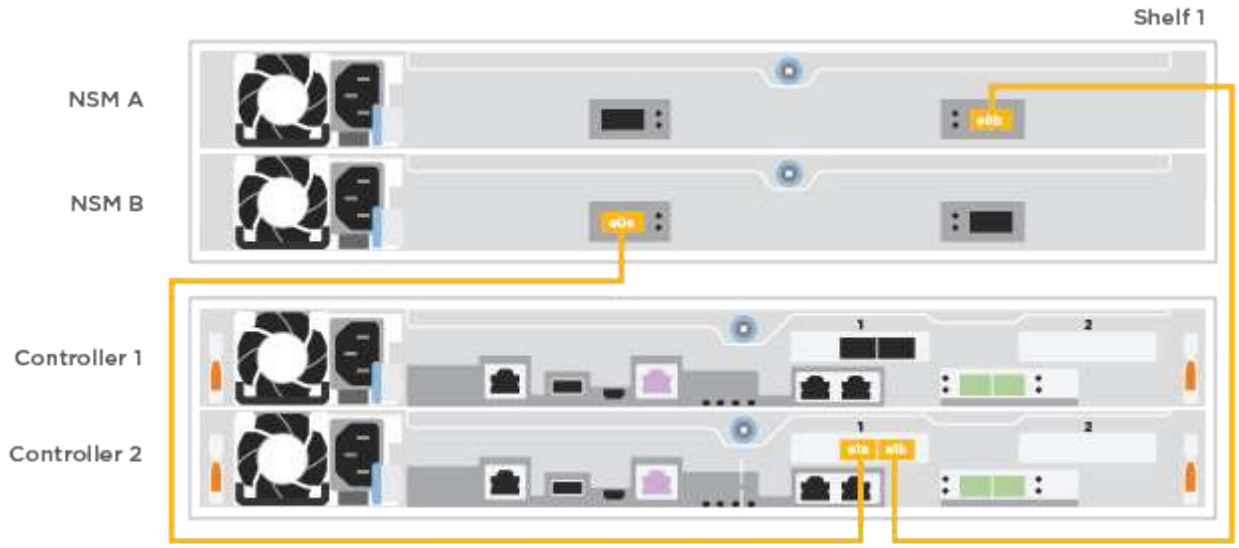

2. Cable controller B to the shelf.

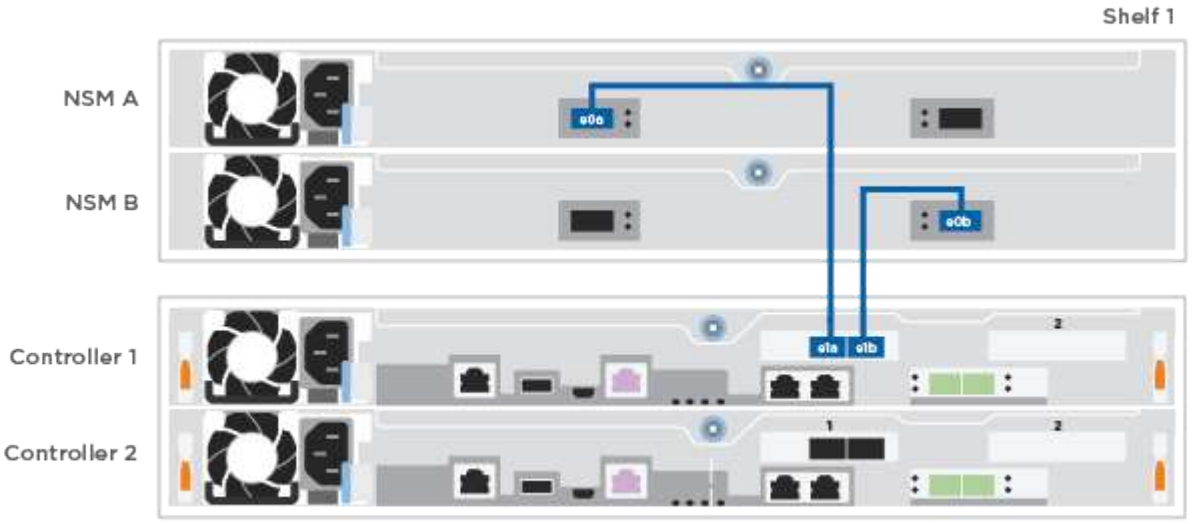

Step 5: Complete system setup

Complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: If network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

2. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

3. Use the animation to connect your laptop to the Management switch:

   Animation - Connect your laptop to the Management switch

4. Select an ONTAP icon listed to discover:

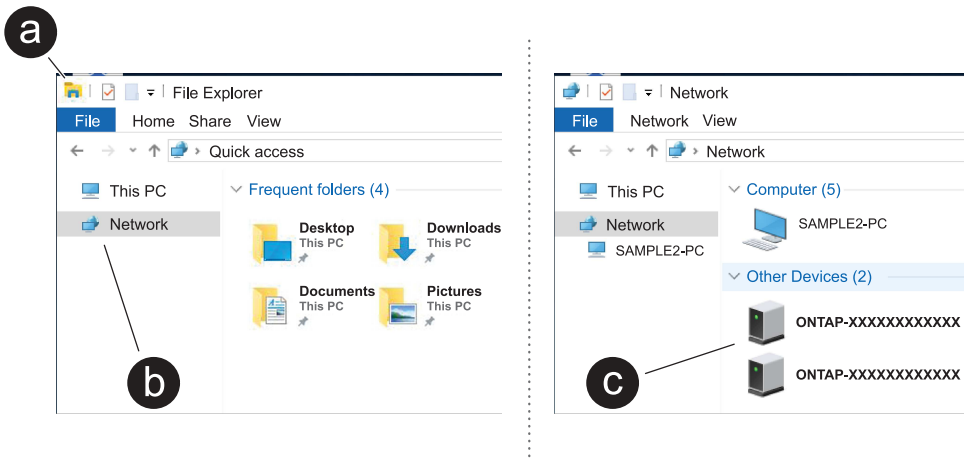

   a. Open File Explorer.

   b. Click **Network** in the left pane and right-click and select **refresh**.

   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

   XXXXX is the system serial number for the target node.

   System Manager opens.

5. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
6. Verify the health of your system by running Config Advisor.
7. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the laptop or console to the switch on the management subnet.

   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.
3. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

4. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the ONTAP Configuration Guide.
5. Verify the health of your system by running Config Advisor.
6. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain AFF A250 hardware

For the AFF A250 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

Fan

The fan cools the controller.

Mezzanine card

A Mezzanine card is a printed circuit board that plugs directly into another plug-in card.

NVEM battery

A battery is included with the controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real-time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF A250

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.
- You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

About this task

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
  - The *impaired* node is the controller on which you are performing maintenance.
  - The *healthy* node is the HA partner of the impaired controller.

Check encryption key support and status - AFF A250

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller - AFF A250

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller displays…	Then…
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: Systems in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *impaired_node_name* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - AFF A250

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller module, you must first remove the controller module from the system, and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Unplug the I/O cables from the controller module.
5. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

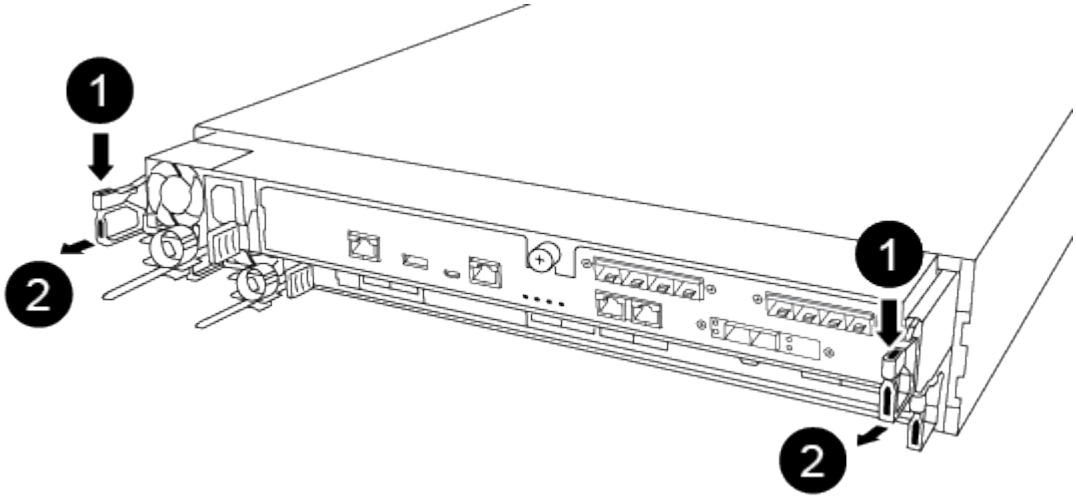

1	Lever
2	Latching mechanism

6. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
7. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

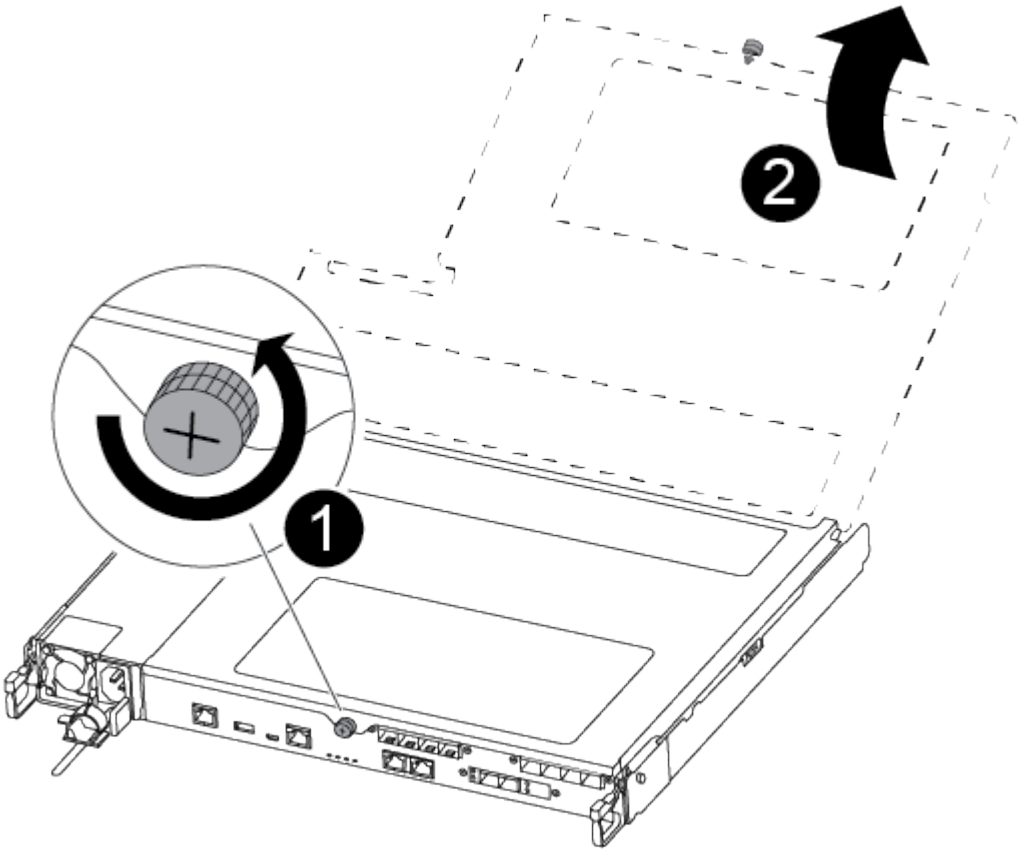

1	Thumbscrew
2	Controller module cover.

8. Lift out the air duct cover.

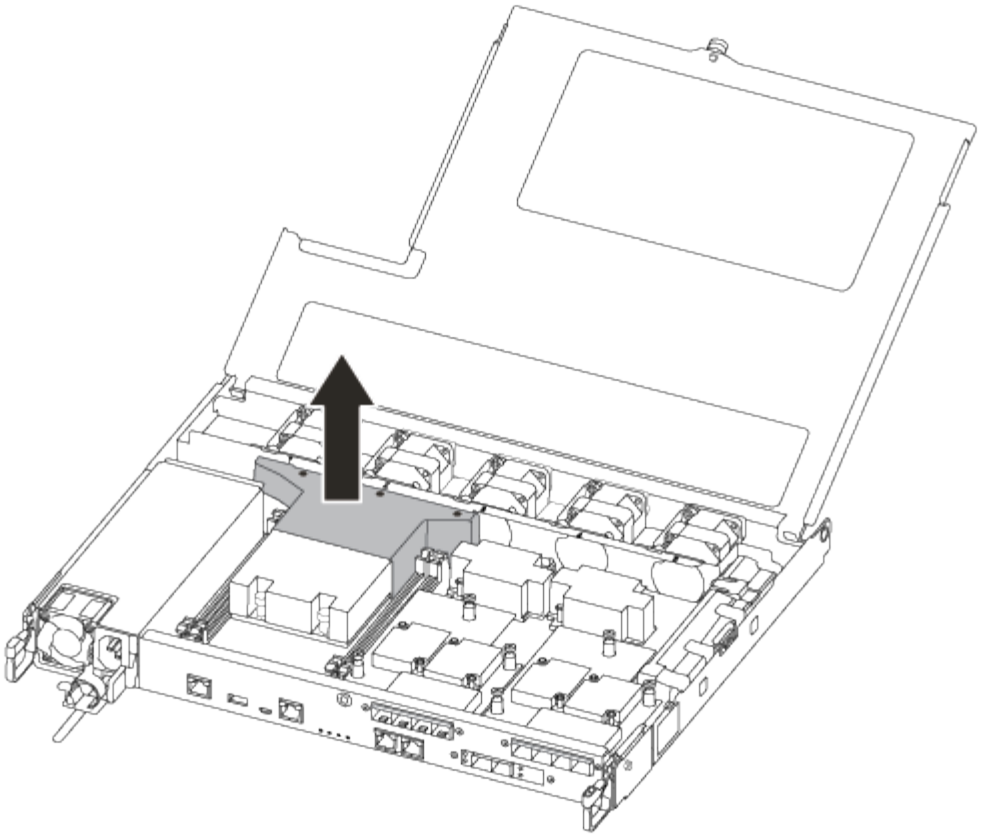

Step 2: Replace the boot media

You locate the failed boot media in the controller module by removing the air duct on the controller module before you can replace the boot media.

You need a #1 magnetic Phillips head screwdriver to remove the screw that holds the boot media in place. Due to the space constraints within the controller module, you should also have a magnet to transfer the screw on to so that you do not lose it.

You can use the following video or the tabulated steps to replace the boot media:

Animation - Replace the boot media

1. Locate and replace the impaired boot media from the controller module.

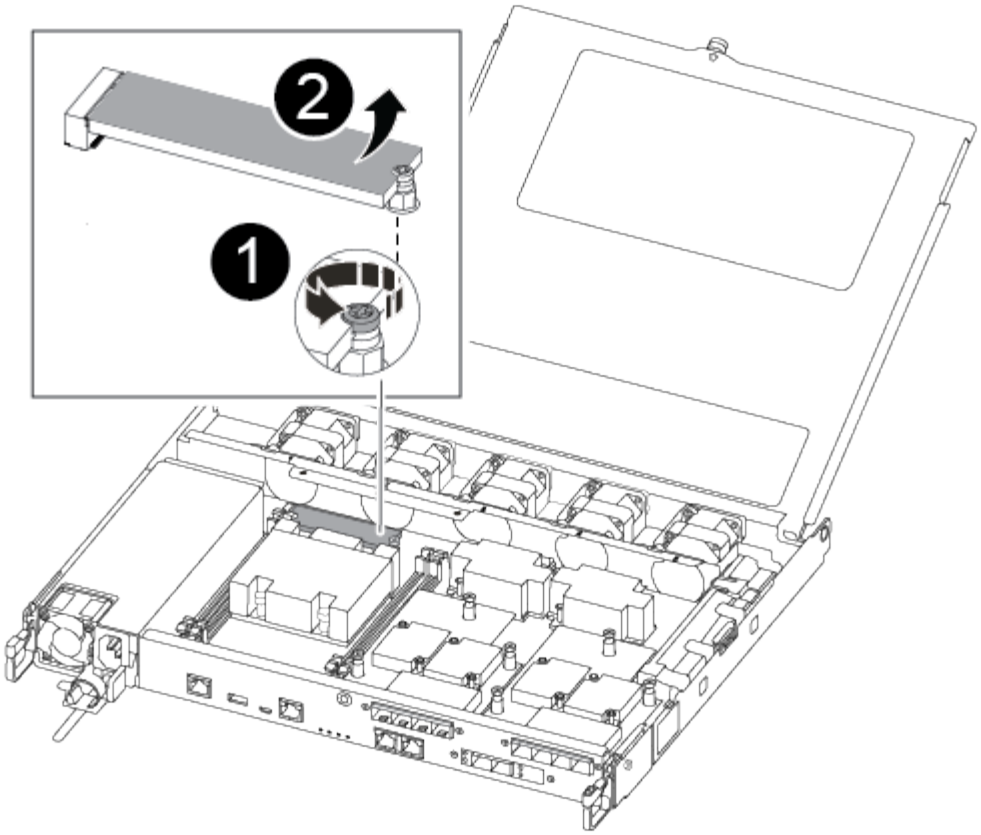

1	Remove the screw securing the boot media to the motherboard in the controller module.
2	Lift the boot media out of the controller module.

2. Using the #1 magnetic screwdriver, remove the screw from the impaired boot media, and set it aside safely on the magnet.
3. Gently lift the impaired boot media directly out of the socket and set it aside.
4. Remove the replacement boot media from the antistatic shipping bag and align it into place on the controller module.
5. Using the #1 magnetic screwdriver, insert and tighten the screw on the boot media.

> Do not apply force when tightening the screw on the boot media; you might crack it.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed is without a boot image so you need to transfer a boot image using a USB flash drive.

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download

button.

- If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
2. Download the service image to your work space on your laptop.
3. Unzip the service image.

   If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

   There are two folders in the unzipped service image file:

   - boot
   - efi

4. Copy the efi folder to the top directory on the USB flash drive.

   If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

   The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

5. Remove the USB flash drive from your laptop.
6. If you have not already done so, install the air duct.

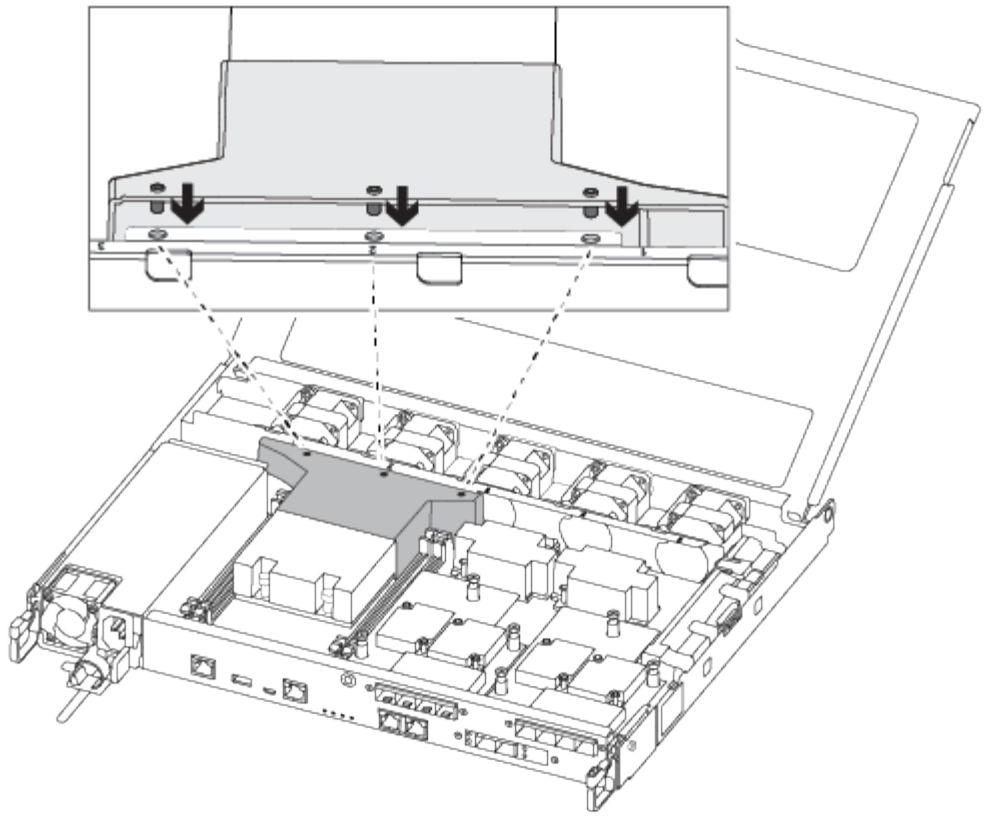

7. Close the controller module cover and tighten the thumbscrew.

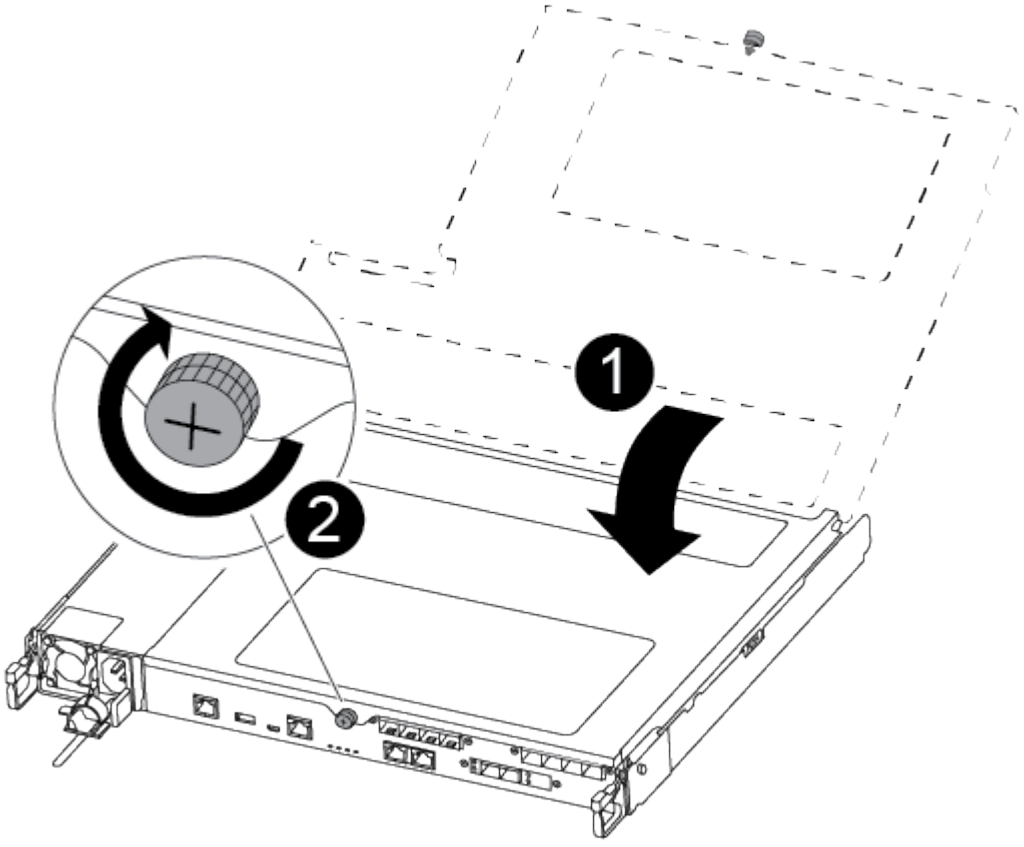

1	Controller module cover
2	Thumbscrew

8. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
9. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

10. Push the controller module all the way into the chassis:
11. Place your index fingers through the finger holes from the inside of the latching mechanism.
12. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
13. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

    The controller module should be fully inserted and flush with the edges of the chassis.

14. Reconnect the controller module I/O cables.
15. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

    The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

16. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

    If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

17. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

    The system begins to boot and stops at the LOADER prompt.

Boot the recovery image - AFF A250

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.

3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.

5. Give back the controller using the `storage failover giveback -fromnode local` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

   NOTE: If the process fails, contact NetApp Support.

Restore encryption - AFF A250

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    `security key-manager key query -restored false.`

    ```
    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    `storage failover giveback -fromnode local`

13. Restore automatic giveback, if you disabled it, by entering the following command.

    `storage failover modify -node local -auto-giveback true`

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    `system node autosupport invoke -node * -type all -message MAINT=END`

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF A250

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF A250

To replace the chassis, you must move the bezel, controller modules, and NVMe drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in its place.

About this task

- All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A250

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Replace hardware - AFF A250

To replace the chassis, you move the power supplies, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis from with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

Use the following video or the tabulated steps to replace the chassis; it assumes the removal and replacement of the bezel:

Animation - Replace the chassis

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

   > If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

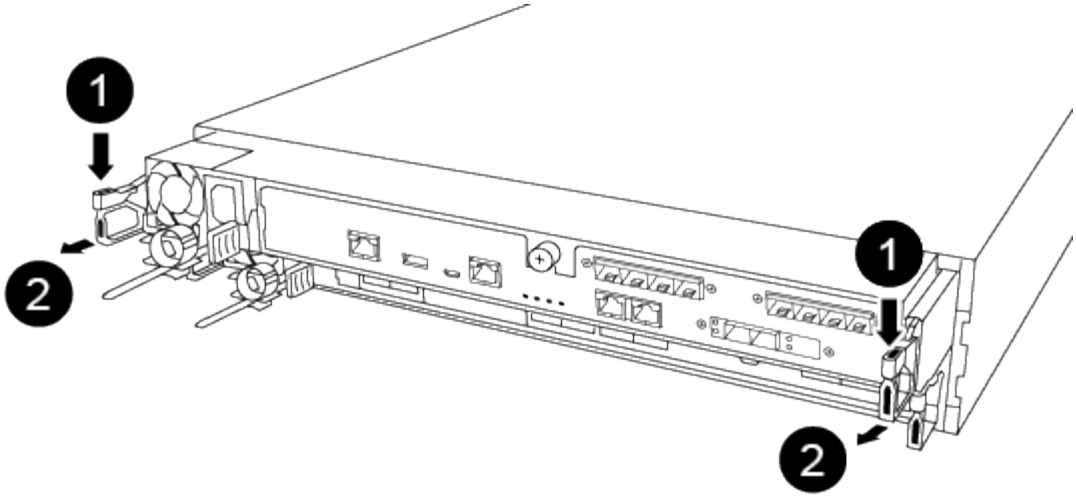

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.
5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up

and against the drive holder.

Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot the system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Insert the controller module into the chassis:

   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

   The controller module should be fully inserted and flush with the edges of the chassis.

   a. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect

the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - AFF A250

You must verify the HA state of the chassis, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis *HA-state*`

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement- AFF A250

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct

recovery procedure to determine whether you should use this procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller module - AFF A250

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - AFF A250

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

Use the following video or the tabulated steps to replace a controller module:

Animation - Replace a controller module

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

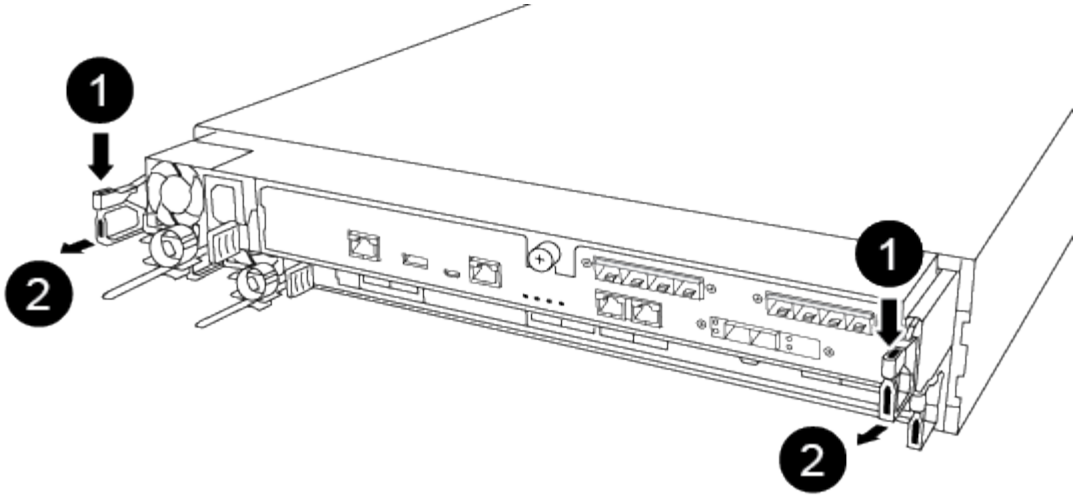

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

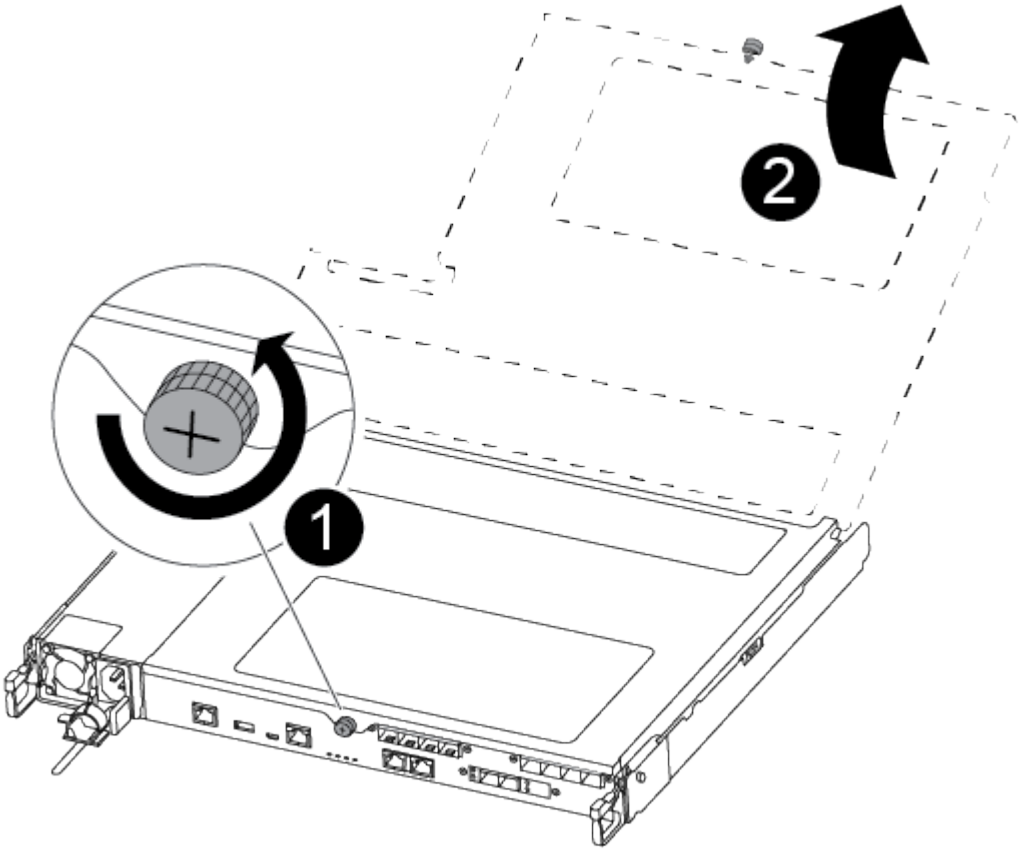

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

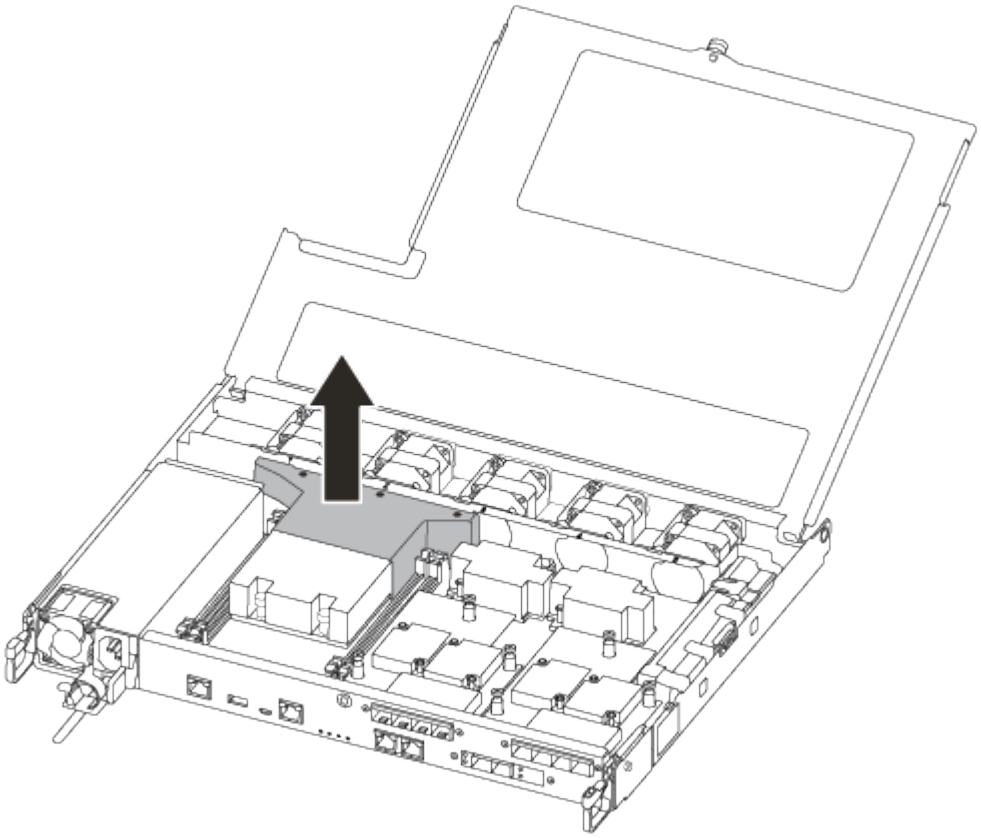

Step 2: Move the power supply

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

1. Disconnect the power supply.
2. Open the power cable retainer, and then unplug the power cable from the power supply.
3. Unplug the power cable from the power source.
4. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

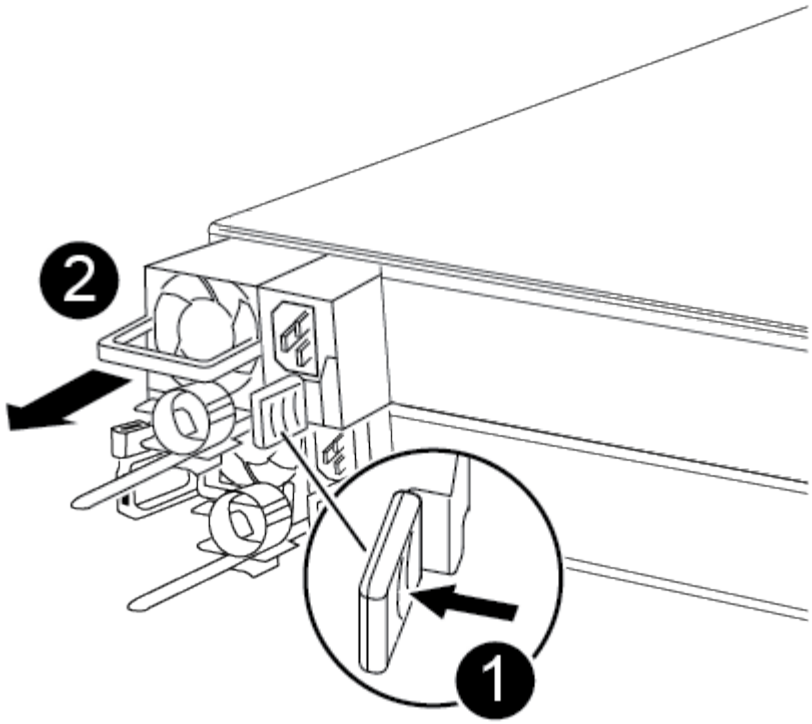

1	Blue power supply locking tab
2	Power supply

5. Move the power supply to the new controller module, and then install it.
6. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

You must move the fans from the impaired controller module to the replacement module when replacing a failed controller module.

1. Remove the fan module by pinching the side of the fan module, and then lifting the fan module straight out of the controller module.

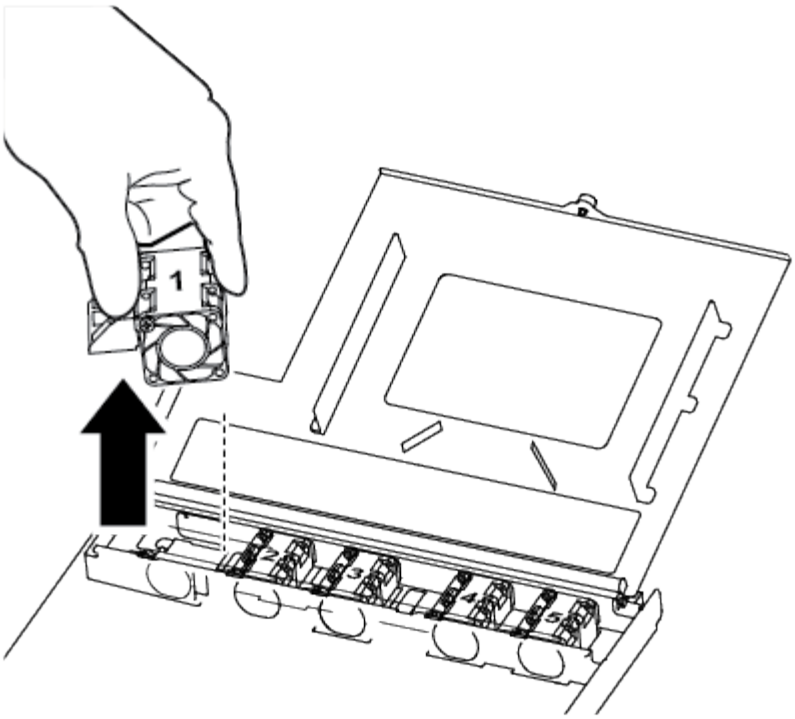

1	Fan module

2. Move the fan module to the replacement controller module, and align the edges of the fan module with the opening in the controller module, and then slide the fan module in.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the boot media

You must move the boot media device from the impaired controller module to the replacement controller module.

You need a #1 magnetic Phillips head screwdriver to remove the screw that holds the boot media in place. Due to the space constraints within the controller module, you should also have a magnet to transfer the screw on to so that you do not lose it.

1. Locate and move the boot media from the impaired controller module to the replacement controller module.

   The boot media is located under the air duct cover you removed earlier in this procedure.

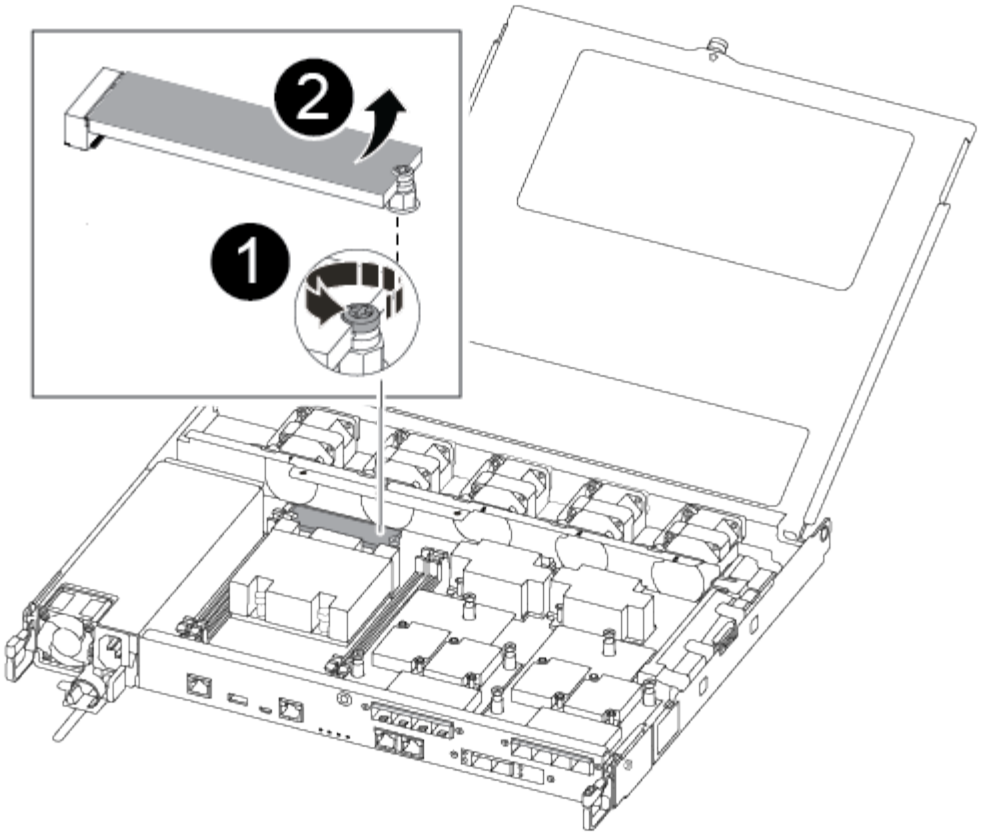

1	Remove the screw securing the boot media to the motherboard in the impaired controller module.
2	Lift the boot media out of the impaired controller module.

2. Using the #1 magnetic screwdriver, remove the screw from the boot media, and set it aside safely on the magnet.
3. Gently lift the boot media directly out of the socket and align it into place in the replacement controller module.
4. Using the #1 magnetic screwdriver, insert and tighten the screw on the boot media.

> Do not apply force when tightening the screw on the boot media; you might crack it.

Step 5: Move the DIMMs

To move the DIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

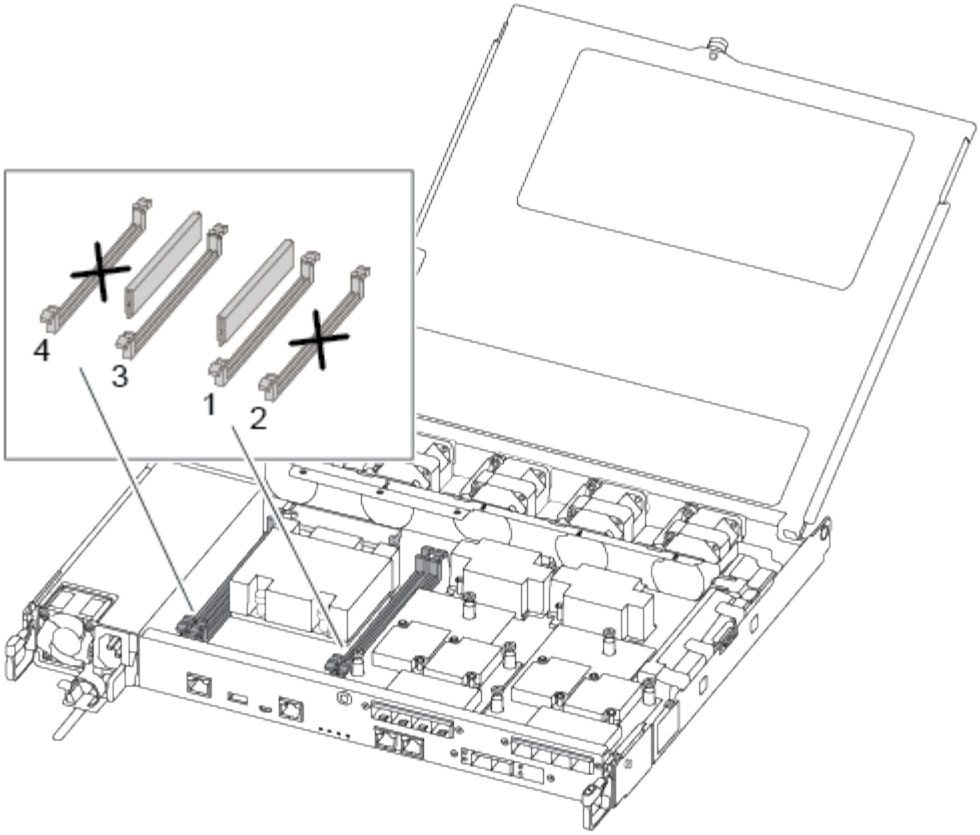

Install each DIMM into the same slot it occupied in the impaired controller module.

1. Slowly push apart the DIMM ejector tabs on either side of the DIMM, and slide the DIMM out of the slot.

   Hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

2. Locate the corresponding DIMM slot on the replacement controller module.
3. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the DIMM squarely into the socket.

   The DIMMs fit tightly in the socket. If not, reinsert the DIMM to realign it with the socket.

4. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.
5. Repeat these steps for the remaining DIMM.

Step 6: Move a mezzanine card

To move a mezzanine card, you must remove the cabling and any QSFPs and SFPs from the ports, move the mezzanine card to the replacement controller, reinstall any QSFPs and SFPs onto the ports, and cable the ports.

1. Locate and move the mezzanine cards from your impaired controller module.

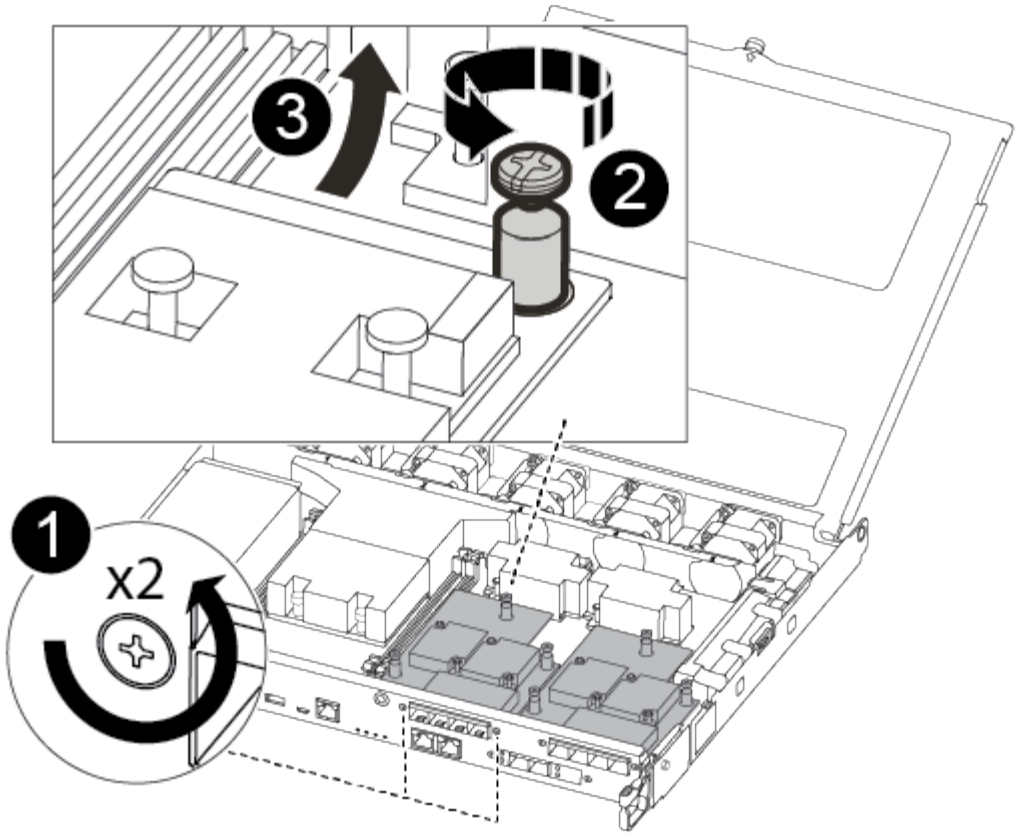

1	Remove screws on the face of the controller module.
2	Loosen the screw in the controller module.
3	Move the mezzanine card.

2. Unplug any cabling associated with the mezzanine card.

   Make sure that you label the cables so that you know where they came from.

   a. Remove any SFP or QSFP modules that might be in the mezzanine card and set it aside.
   b. Using the #1 magnetic screwdriver, remove the screws from the face of the impaired controller module and from the mezzanine card, and set them aside safely on the magnet.
   c. Gently lift the mezzanine card out of the socket and move it to the same position in the replacement controller.
   d. Gently align the mezzanine card into place in the replacement controller.
   e. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the replacement controller module and on the mezzanine card.

      Do not apply force when tightening the screw on the mezzanine card; you might crack it.

3. Repeat these steps if there is another mezzanine card in the impaired controller module.

4. Insert the SFP or QSFP modules that were removed onto the mezzanine card.

Step 7: Move the NV battery

When replacing the controller module, you must move the NV battery from the impaired controller module to the replacement controller module.

1. Locate and move the NVMEM battery from your impaired controller module to the replacement controller module.

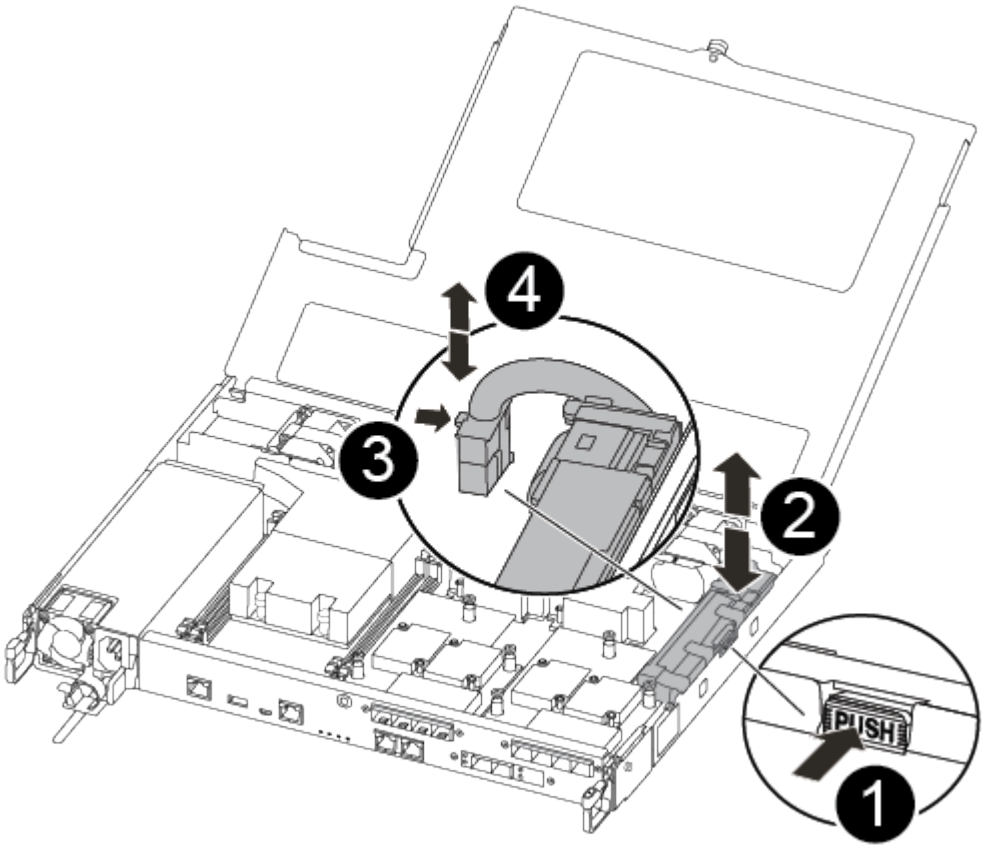

1	Squeeze the clip on the face of the battery plug.
2	Unplug the battery cable from the socket.
3	Grasp the battery and press the blue locking tab marked PUSH.
4	Lift the battery out of the holder and controller module.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.

4. Locate the corresponding NV battery holder on the replacement controller module and align the NV battery to the battery holder.
5. Insert the NV battery plug into the socket.
6. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
7. Press firmly down on the battery pack to make sure that it is locked into place.

Step 8: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

You can use the following illustrations or the written steps to install the replacement controller module in the chassis.

1. If you have not already done so, install the air duct.

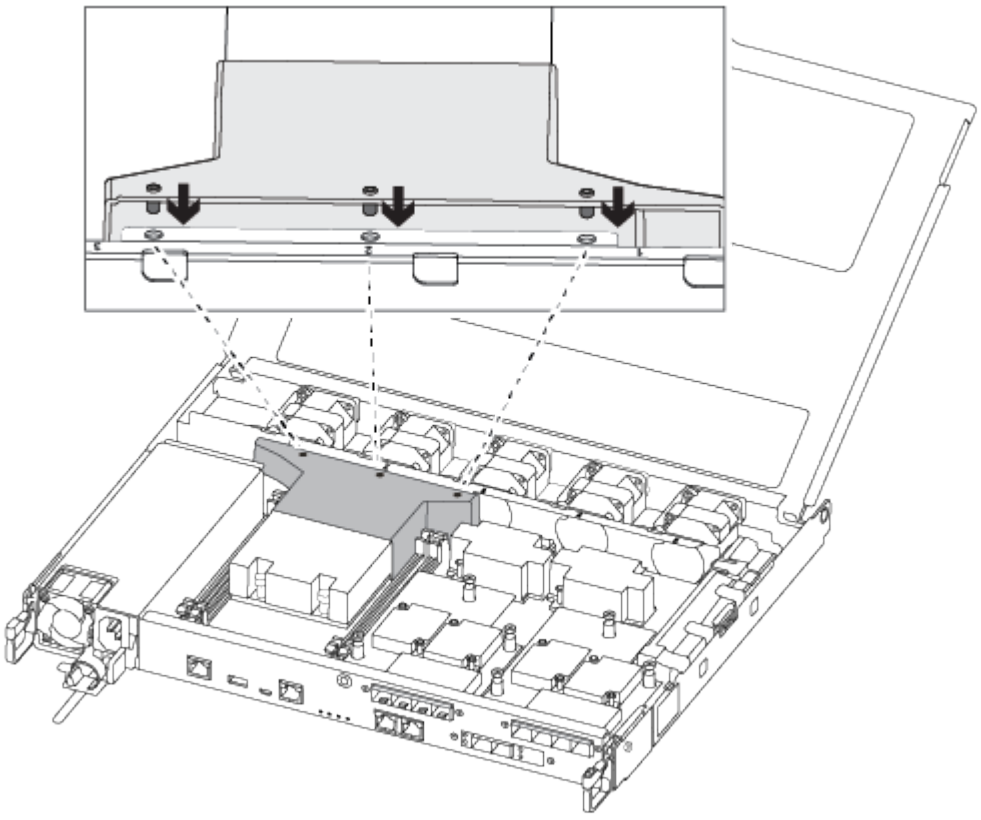

2. Close the controller module cover and tighten the thumbscrew.

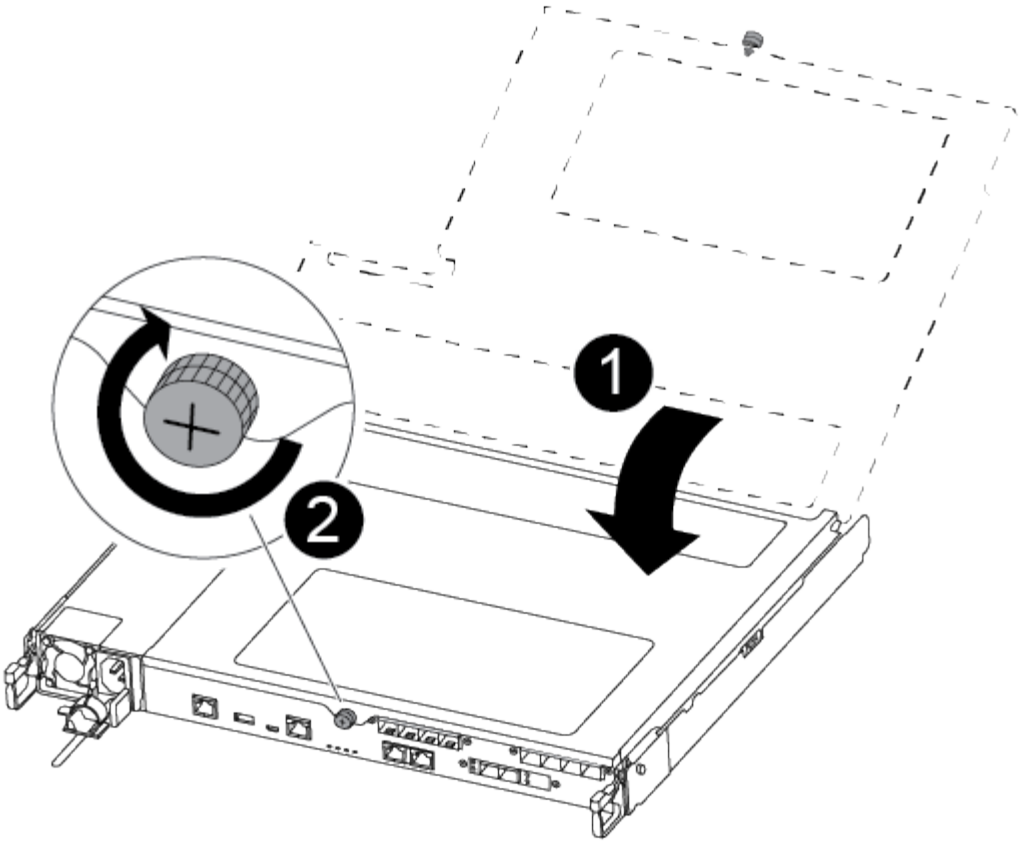

1	Controller module cover
2	Thumbscrew

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching

mechanisms snap into place.

The controller module should be fully inserted and flush with the edges of the chassis.

f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

Restore and verify the system configuration - AFF A250

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - ha
   - mcc
   - mccip
   - non-ha

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF A250

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and

then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF A250

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A250

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

About this task

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

> If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

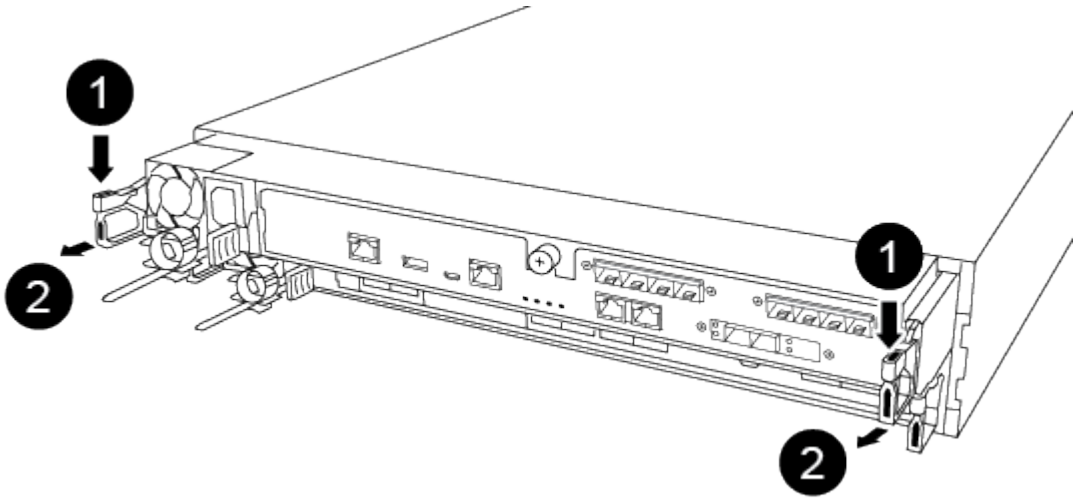

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

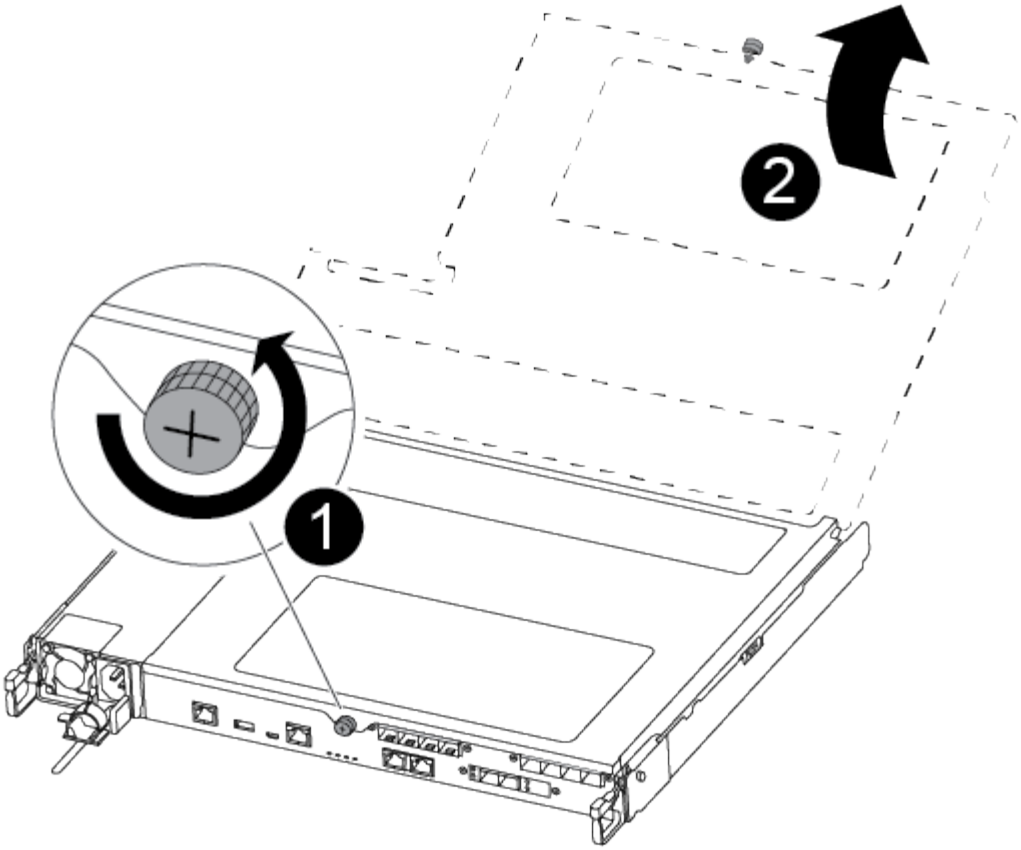

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

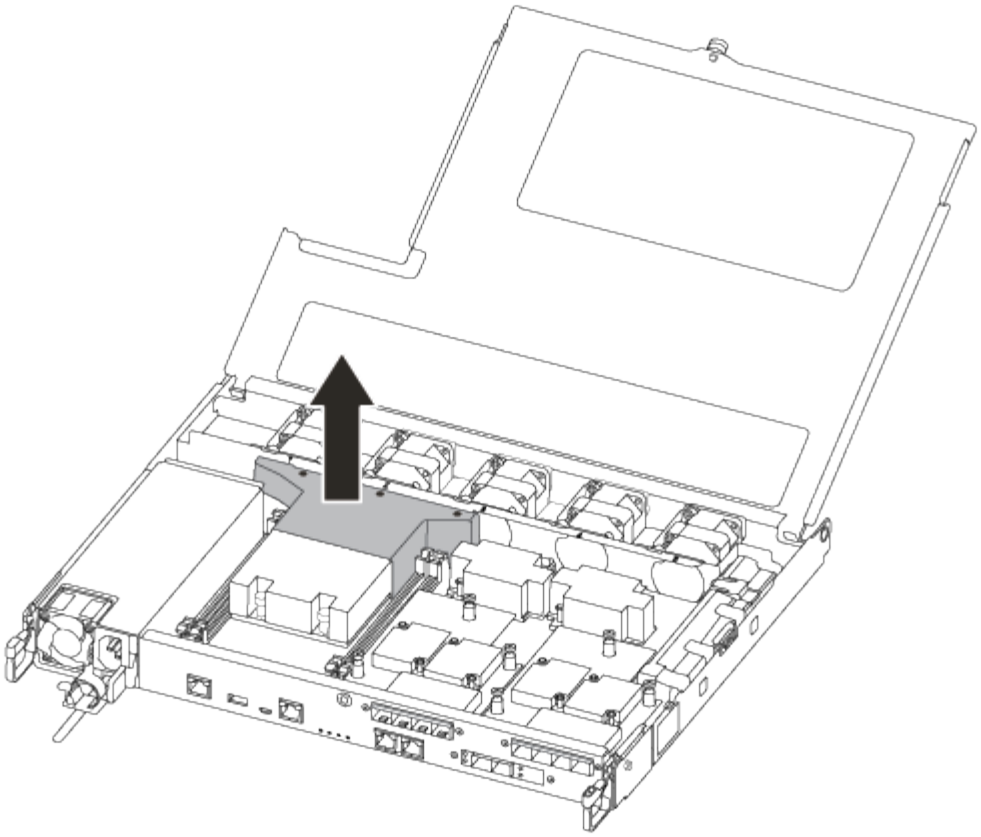

Step 3: Replace a DIMM

To replace a DIMM, you must locate it in the controller module using the DIMM map label on top of the air duct and then replace it following the specific sequence of steps.

Use the following video or the tabulated steps to replace a DIMM:

Animation - Replace a DIMM

1. Replace the impaired DIMM on your controller module.

   The DIMMs are in slot 3 or 1 on the motherboard. Slot 2 and 4 are left empty. Do not attempt to install DIMMs into these slots.

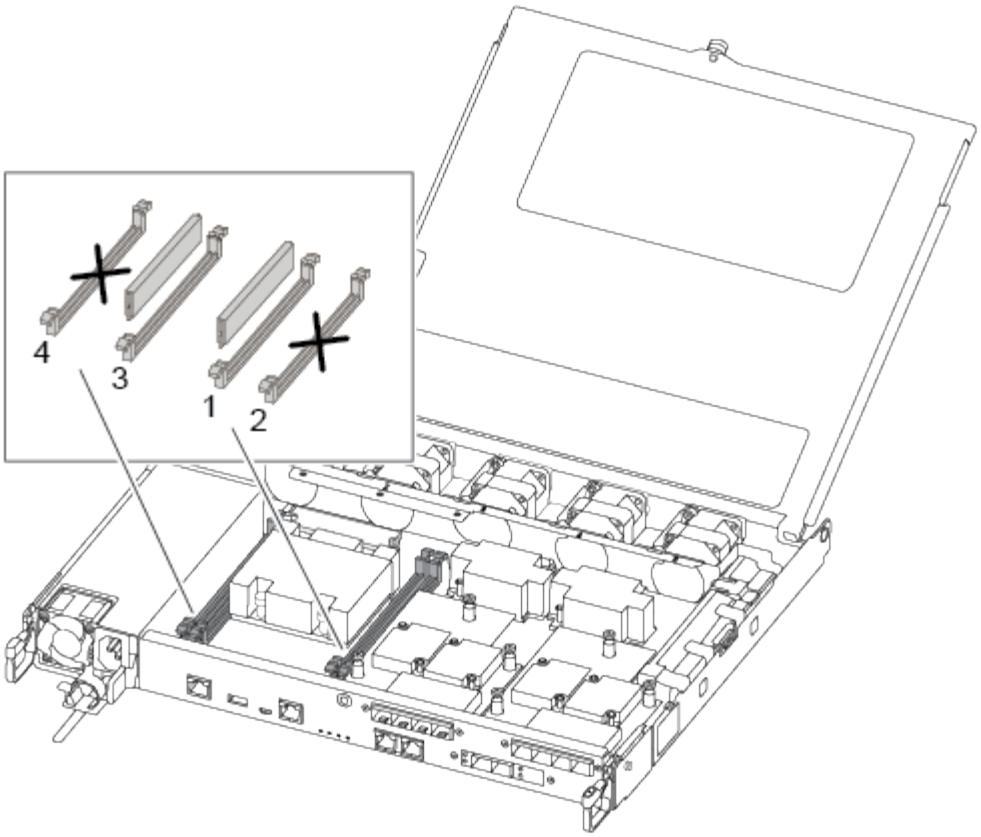

2. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
3. Slowly push apart the DIMM ejector tabs on either side of the DIMM, and slide the DIMM out of the slot.
4. Leave DIMM ejector tabs on the connector in the open position.
5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   Hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

6. Insert the replacement DIMM squarely into the slot.

   The DIMMs fit tightly in the socket. If not, reinsert the DIMM to realign it with the socket.

7. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following illustrations or the written steps to install the replacement controller module in the chassis.

1. If you have not already done so, install the air duct.

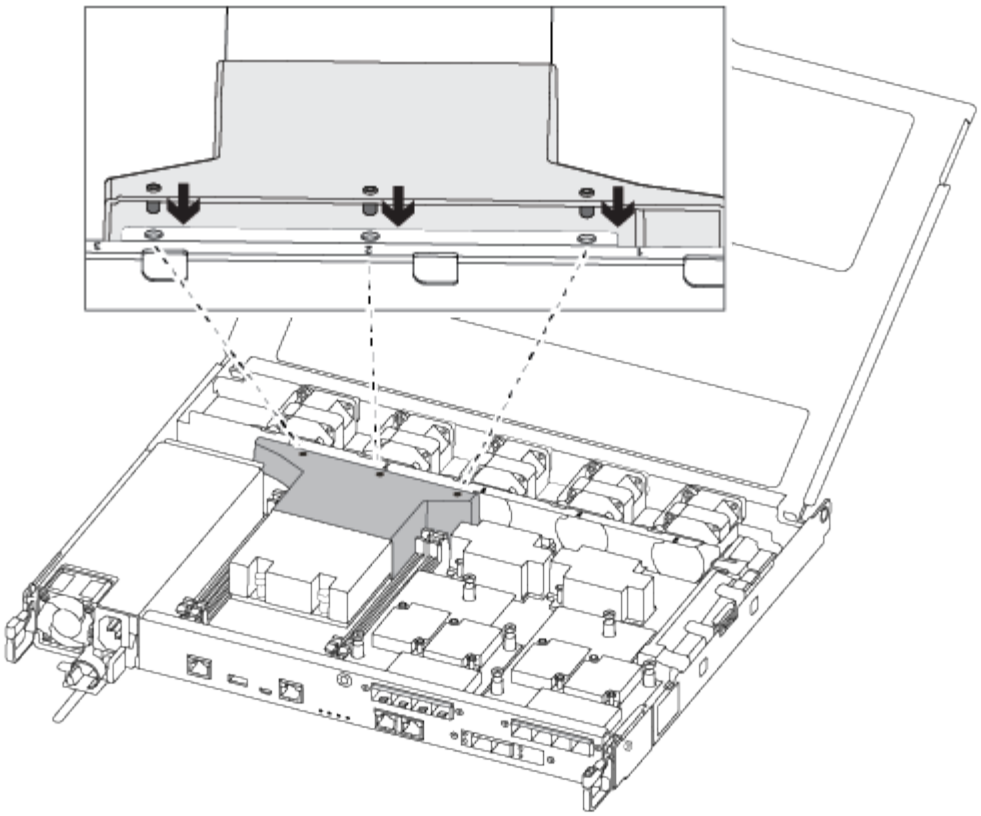

2. Close the controller module cover and tighten the thumbscrew.

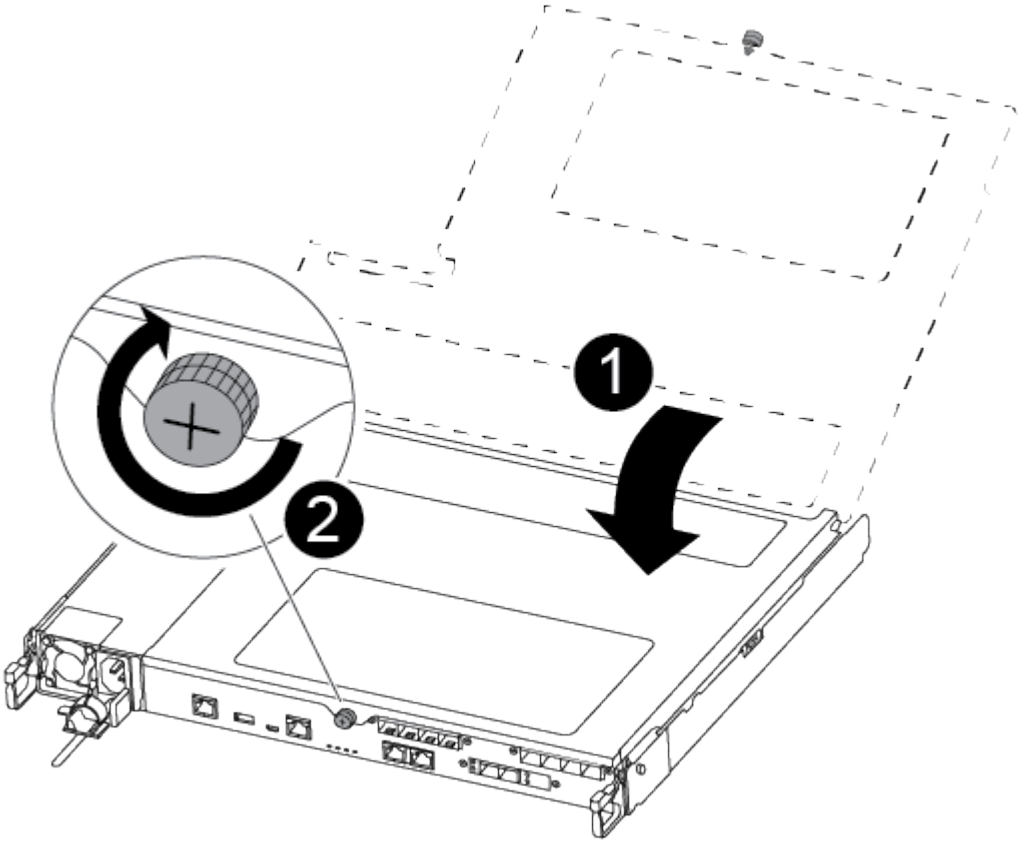

1	Controller module cover
2	Thumbscrew

3. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.
4. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.
5. Recable the system, as needed.
6. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
7. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - AFF A250

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system

console.

The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   > The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan - AFF A250

To replace a fan, remove the failed fan module and replace it with a new fan module.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

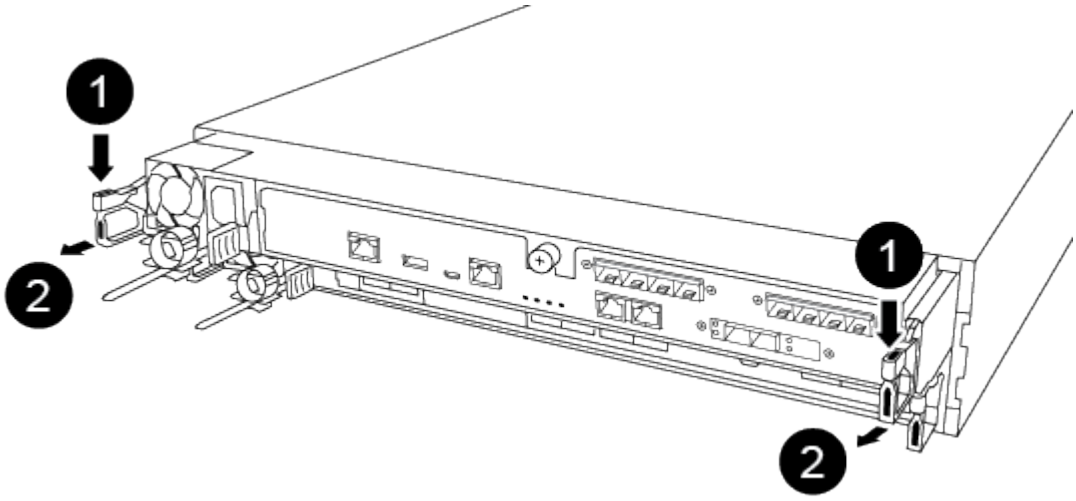

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

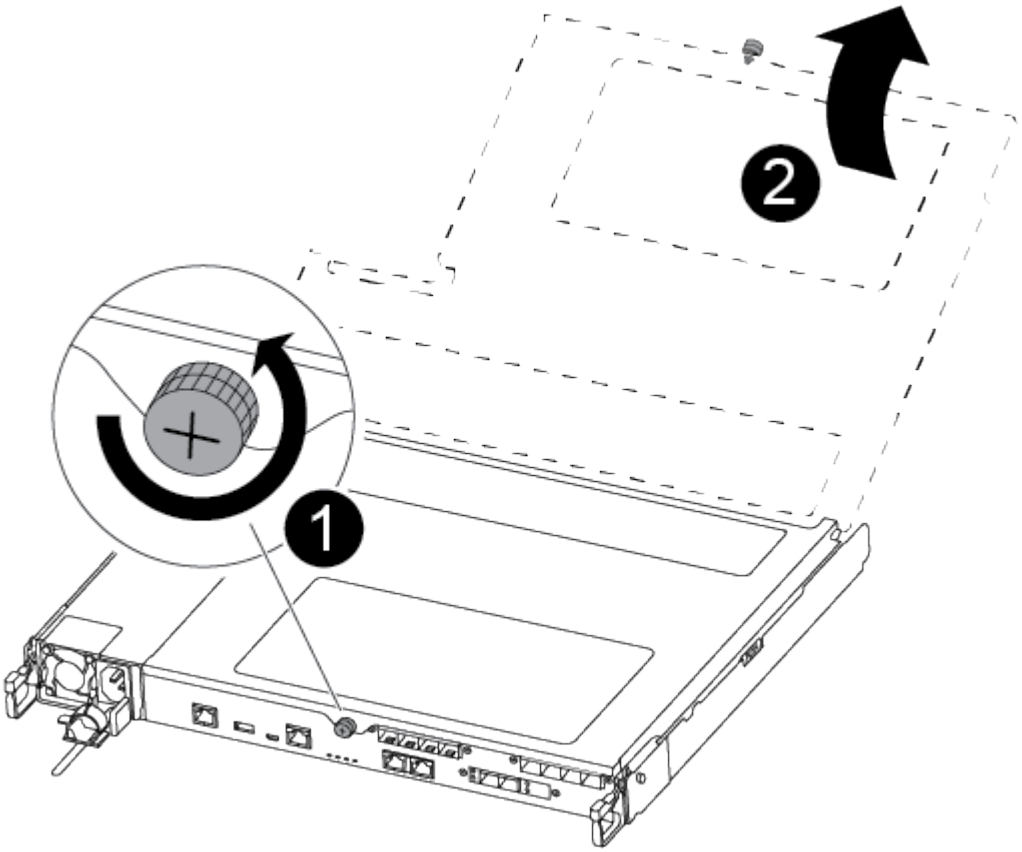

1	Thumbscrew
2	Controller module cover

Step 3: Replace a fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

Use the following video or the tabulated steps to replace a fan:

Animation - Replace a fan

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
2. Remove the fan module by pinching the side of the fan module, and then lifting the fan module straight out of the controller module.

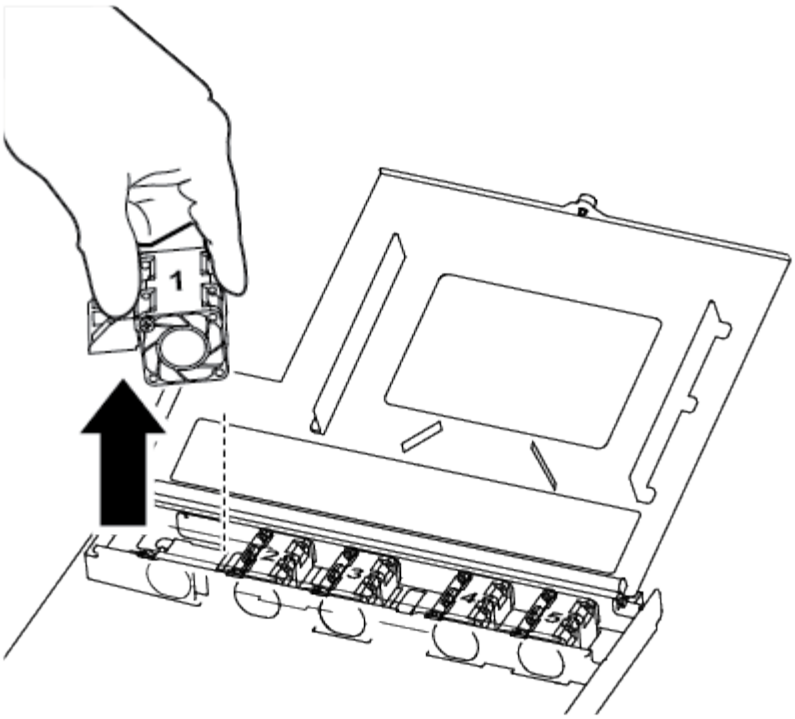

1	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Close the controller module cover and tighten the thumbscrew.

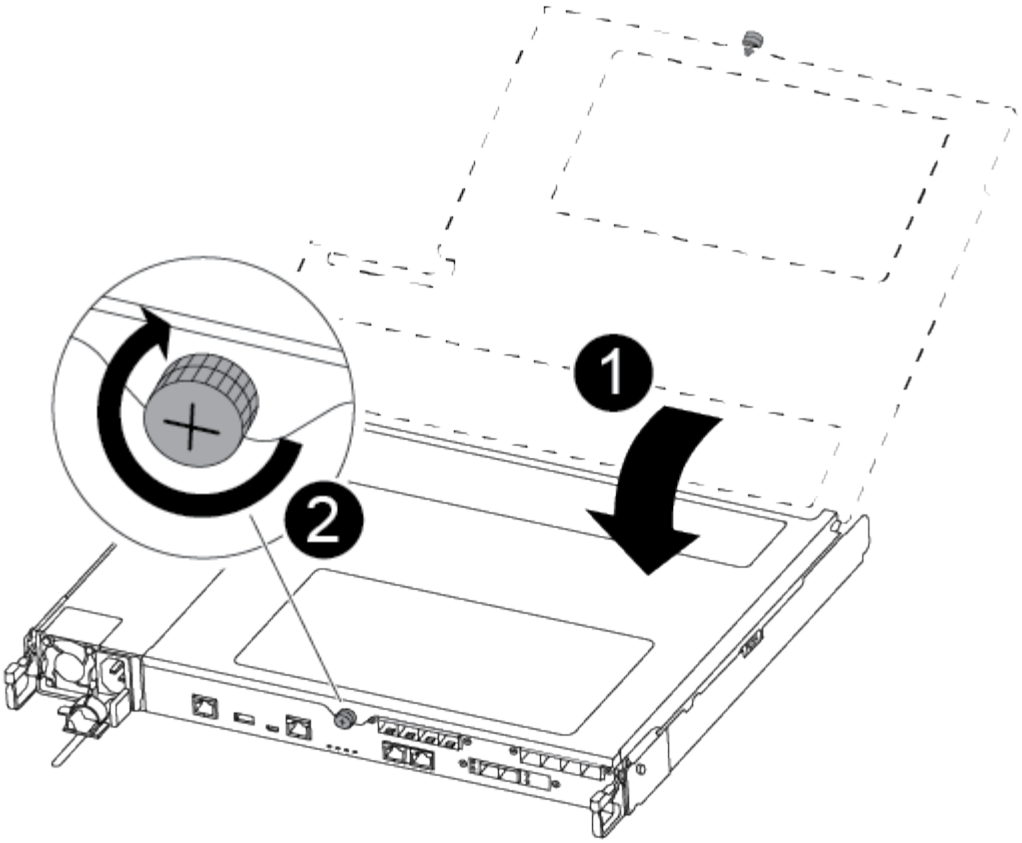

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.

   f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace or install a mezzanine card - AFF A250

To replace a failed mezzanine card, you must remove the cables and any SFP or QSFP modules, replace the card, reinstall the SFP or QSFP modules and recable the cards. To install a new mezzanine card, you must have the appropriate cables and SFP or QSFP modules.

About this task

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

Remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

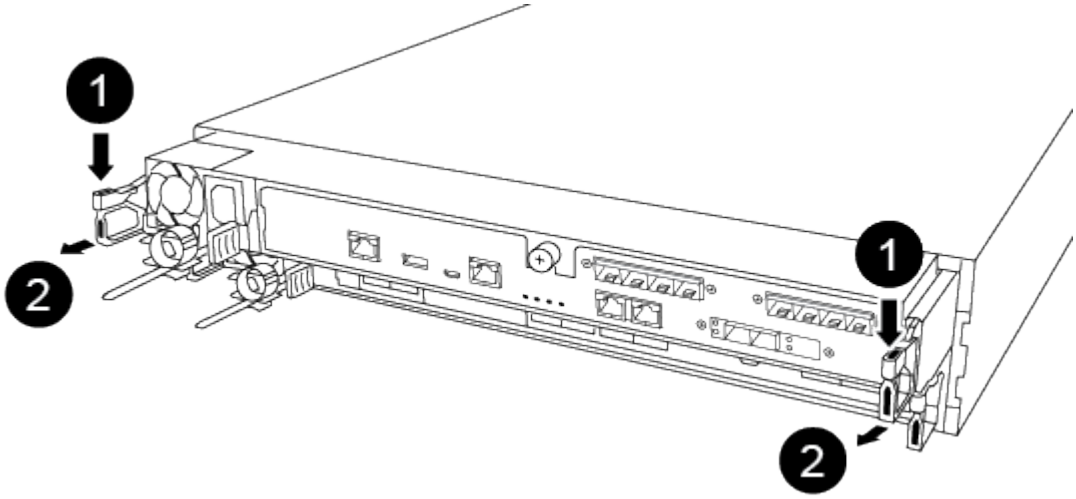

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

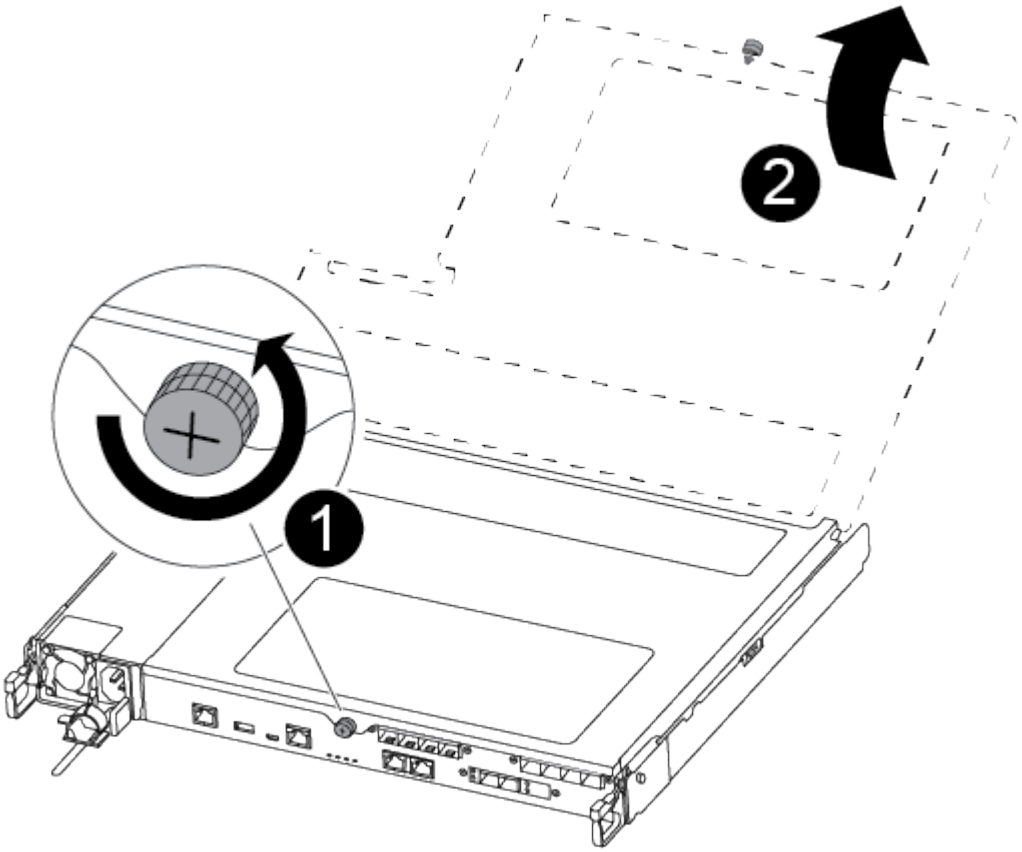

1	Thumbscrew
2	Controller module cover.

Step 3: Replace or install a mezzanine card

To replace a mezzanine card, you must remove the impaired card and install the replacement card; to install a mezzanine card, you must remove the faceplate and install the new card.

Use the following video or the tabulated steps to replace a mezzanine card:

Animation - Replace a mezzanine card

1. To replace a mezzanine card:
2. Locate and replace the impaired mezzanine card on your controller module.

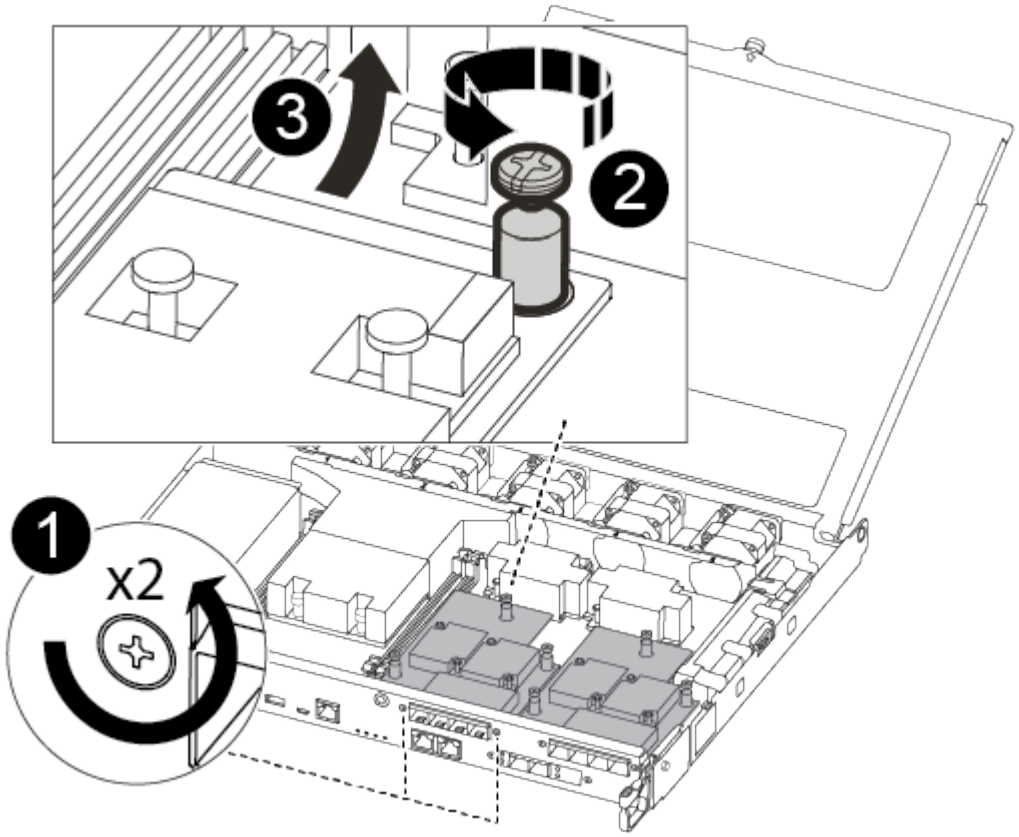

1	Remove screws on the face of the controller module.
2	Loosen the screw in the controller module.
3	Remove the mezzanine card.

a. Unplug any cabling associated with the impaired mezzanine card.

   Make sure that you label the cables so that you know where they came from.

b. Remove any SFP or QSFP modules that might be in the impaired mezzanine card and set it aside.
c. Using the #1 magnetic screwdriver, remove the screws from the face of the controller module and set them aside safely on the magnet.
d. Using the #1 magnetic screwdriver, loosen the screw on the impaired mezzanine card.
e. Using the #1 magnetic screwdriver, gently lift the impaired mezzanine card directly out of the socket and set it aside.
f. Remove the replacement mezzanine card from the antistatic shipping bag and align it to the inside face of the controller module.
g. Gently align the replacement mezzanine card into place.

h. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the controller module and on the mezzanine card.

> Do not apply force when tightening the screw on the mezzanine card; you might crack it.

i. Insert any SFP or QSFP modules that were removed from the impaired mezzanine card to the replacement mezzanine card.

3. To install a mezzanine card:
4. You install a new mezzanine card if your system does not have one.
   a. Using the #1 magnetic screwdriver, remove the screws from the face of the controller module and the faceplate covering the mezzanine card slot, and set them aside safely on the magnet.
   b. Remove the mezzanine card from the antistatic shipping bag and align it to the inside face of the controller module.
   c. Gently align the mezzanine card into place.
   d. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the controller module and on the mezzanine card.

> Do not apply force when tightening the screw on the mezzanine card; you might crack it.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Close the controller module cover and tighten the thumbscrew.

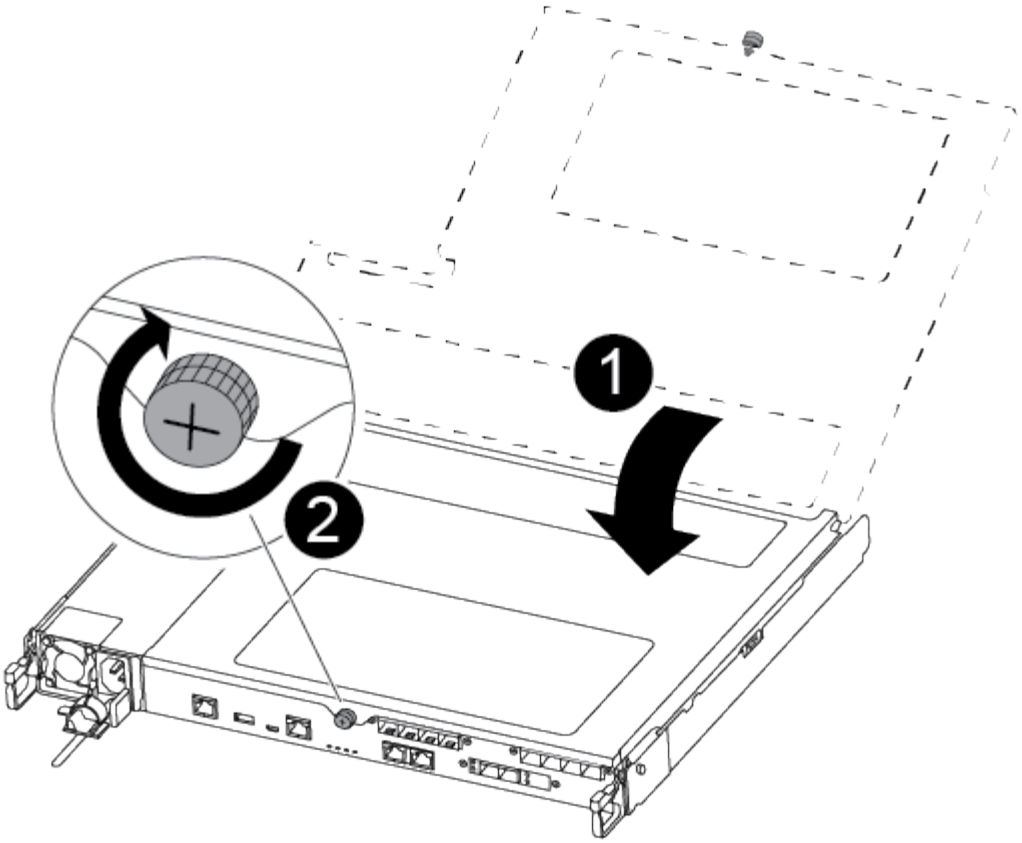

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.

   f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVMEM battery - AFF A250

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

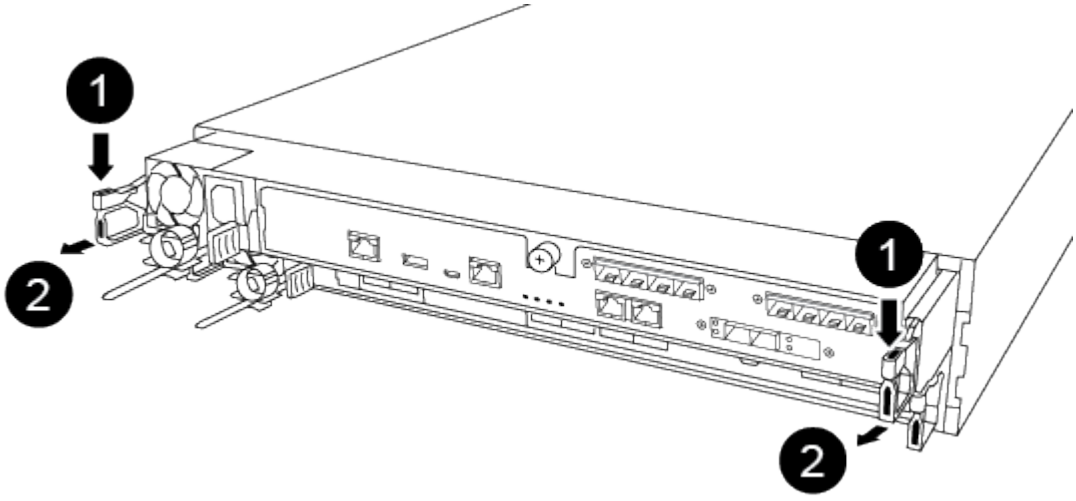

1	Lever

2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

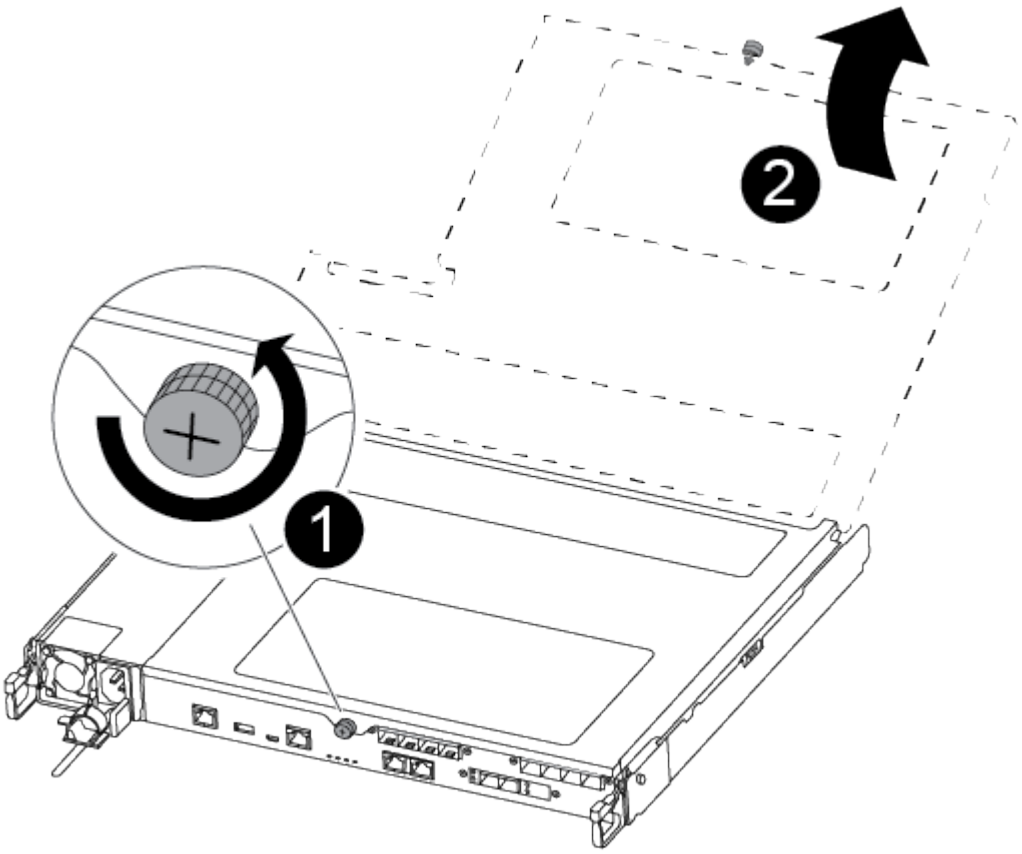

1	Thumbscrew
2	Controller module cover.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module.

Use the following video or the tabulated steps to replace the NVMEM battery:

Animation - Replace the NVMEM battery

1. Locate and replace the impaired NVMEM battery on your controller module.

It is recommended that you follow the illustrated instructions in the order listed.

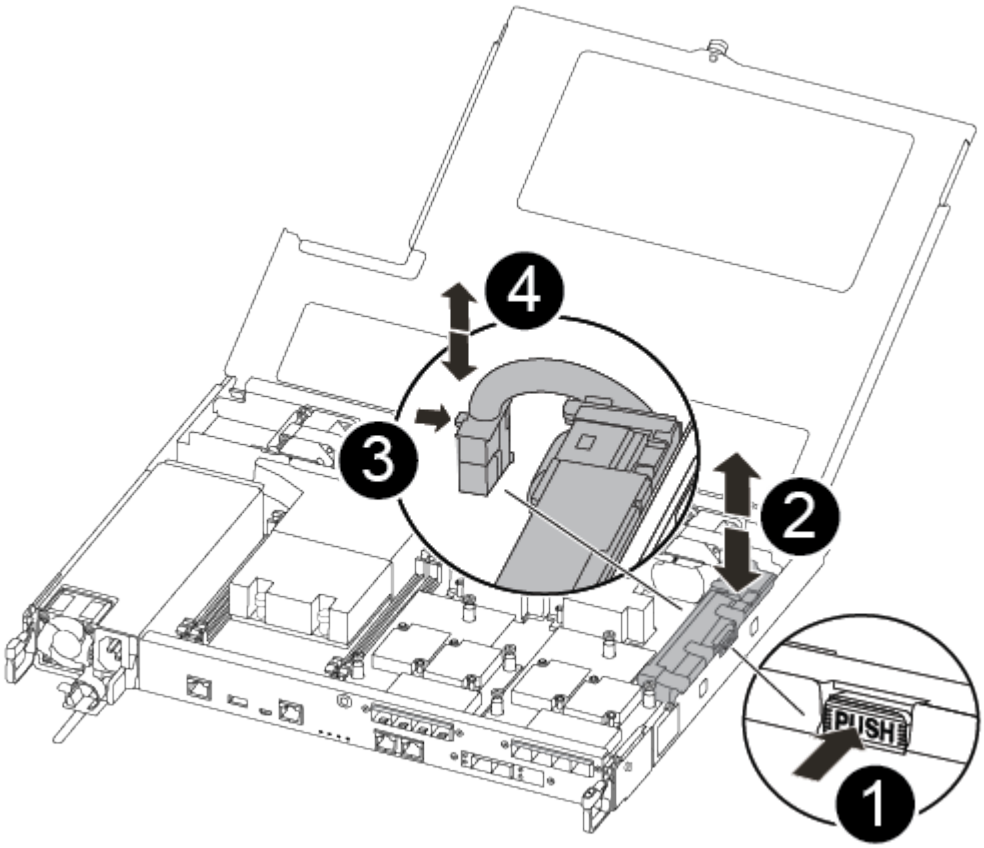

1	Squeeze the clip on the face of the battery plug.
2	Unplug the battery cable from the socket.
3	Grasp the battery and press the blue locking tab marked PUSH.
4	Lift the battery out of the holder and controller module.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module and set it aside.
4. Remove the replacement NV battery from the antistatic shipping bag and align it to the battery holder.
5. Insert the replacement NV battery plug into the socket.
6. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
7. Press firmly down on the battery pack to make sure that it is locked into place.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following illustration or the written steps to install the replacement controller module in the chassis.

1. Close the controller module cover and tighten the thumbscrew.

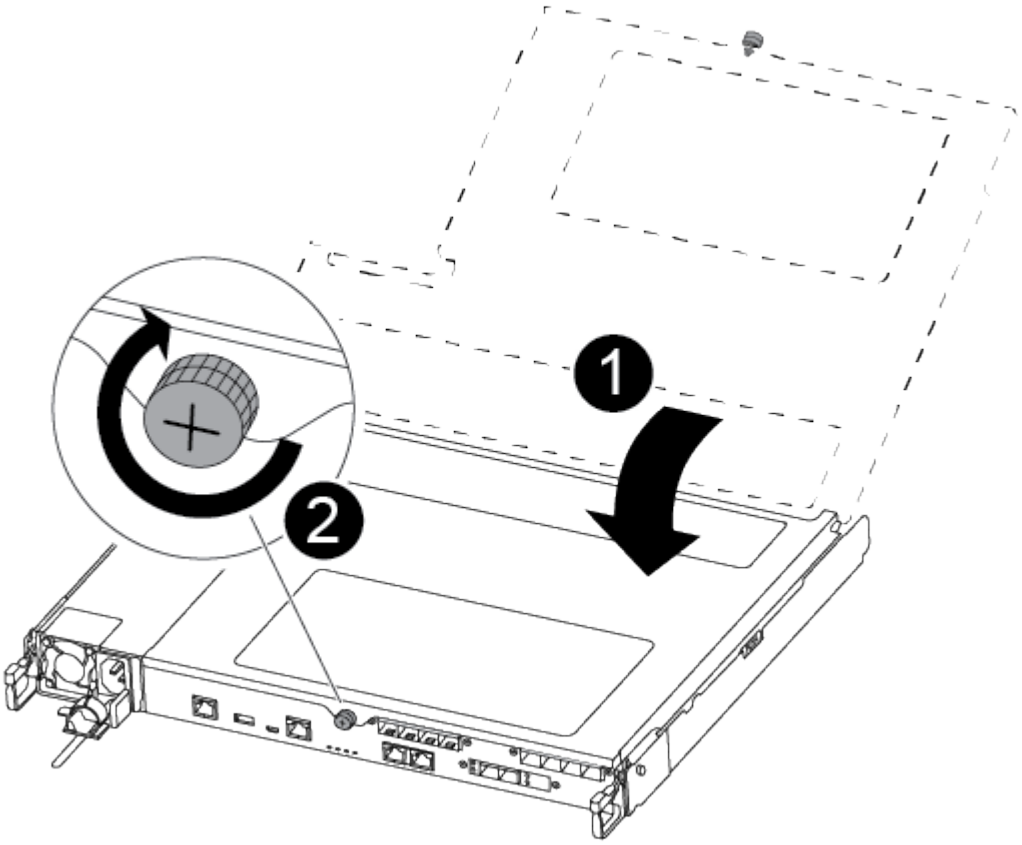

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

The controller module should be fully inserted and flush with the edges of the chassis.

f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.
4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - AFF A250

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one PSU at a time.

> It is a best practice to replace the PSU within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded PSU until the PSU is replaced.

- Power supplies are auto-ranging.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

Use the appropriate procedure for your type of PSU; AC or DC.

Option 1: Replace an AC PSU

Use the following video or the tabulated steps to replace the PSU:

Animation - Replace the AC PSU

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

> The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

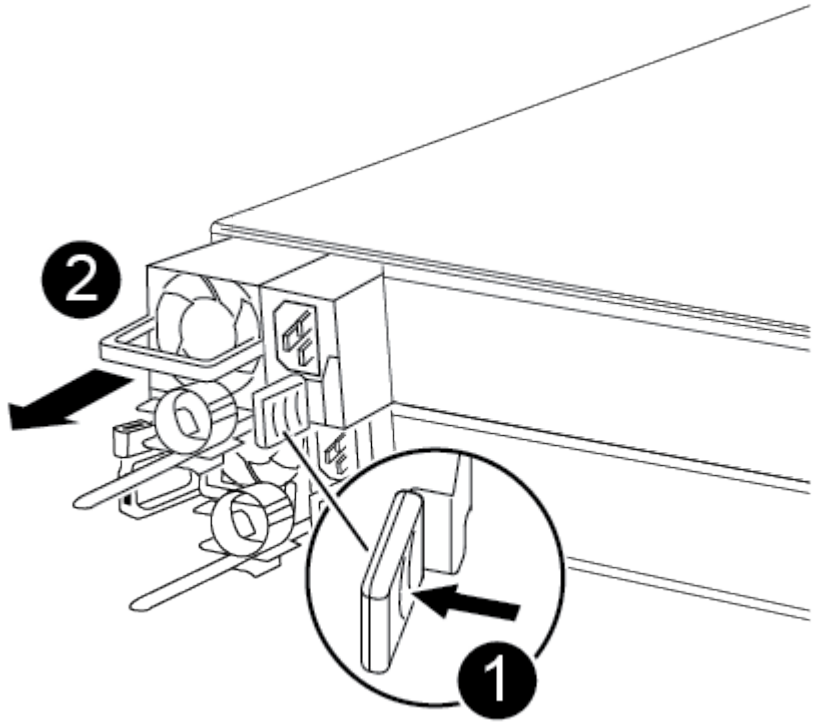

1	Blue PSU locking tab
2	Power supply

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller module.

b. Gently push the PSU into the controller module until the locking tab clicks into place.

The power supplies will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:

a. Reconnect the power cable to the PSU.

b. Secure the power cable to the PSU using the power cable retainer.

Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC power cable connector using the thumb screws on the plug.
   b. Unplug the power cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

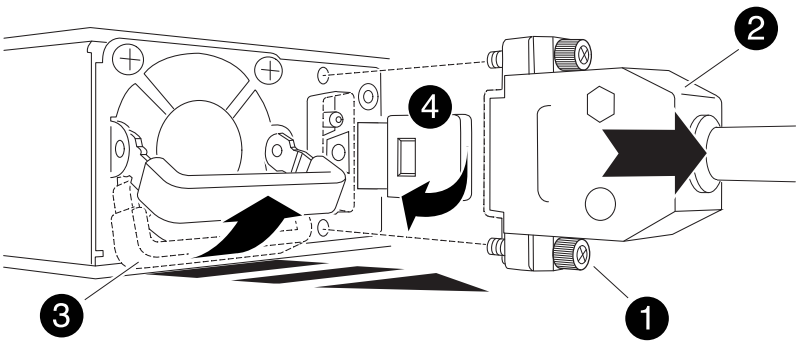

1	Thumb screws
2	D-SUB DC power cable connector

3	Power supply handle
4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A250

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

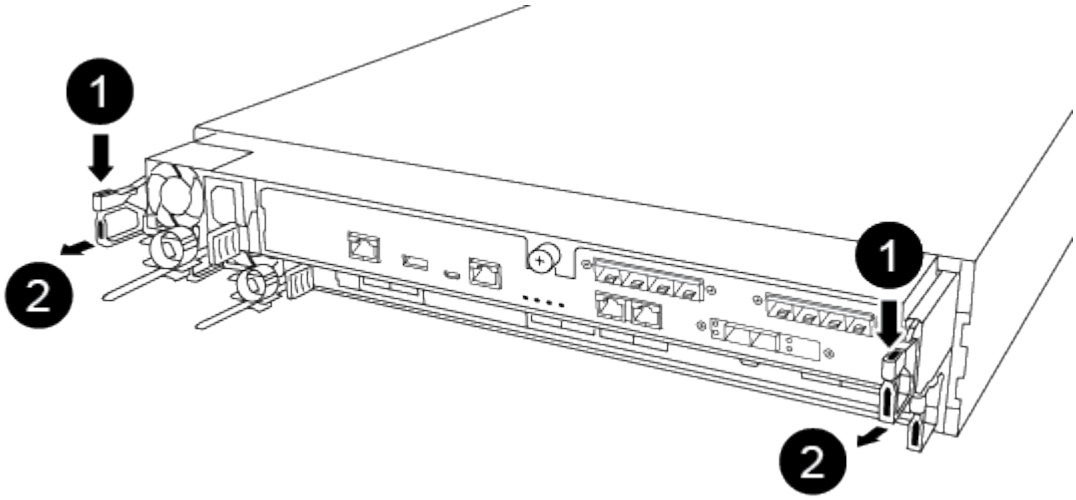

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

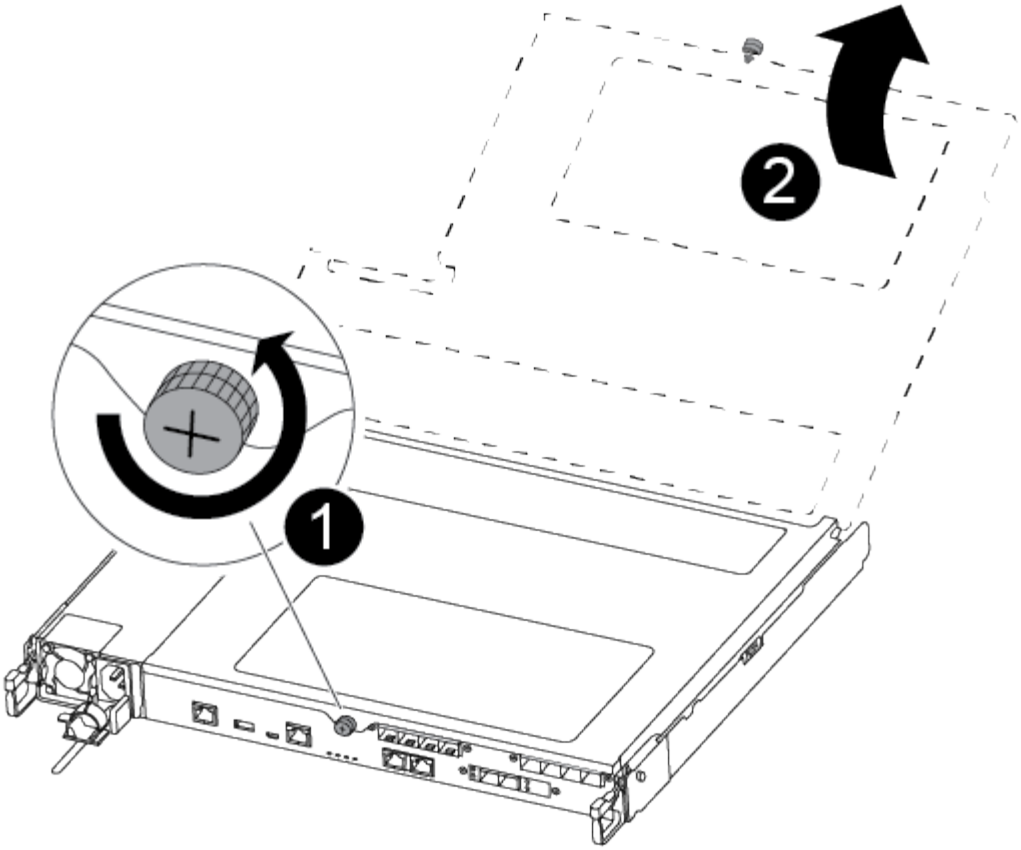

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

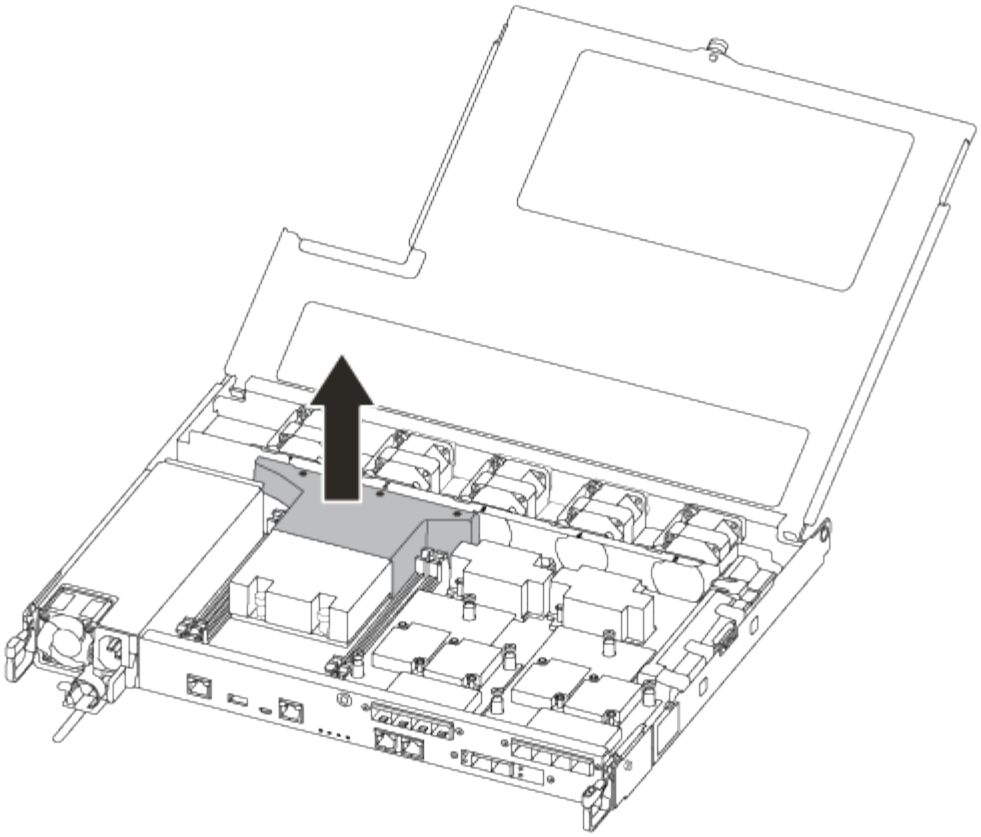

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

Use the following video or the tabulated steps to replace the RTC battery:

Animation - Replace the RTC battery

1. Locate the RTC battery between the heatsink and the midplane and remove it exactly as shown in the graphic.

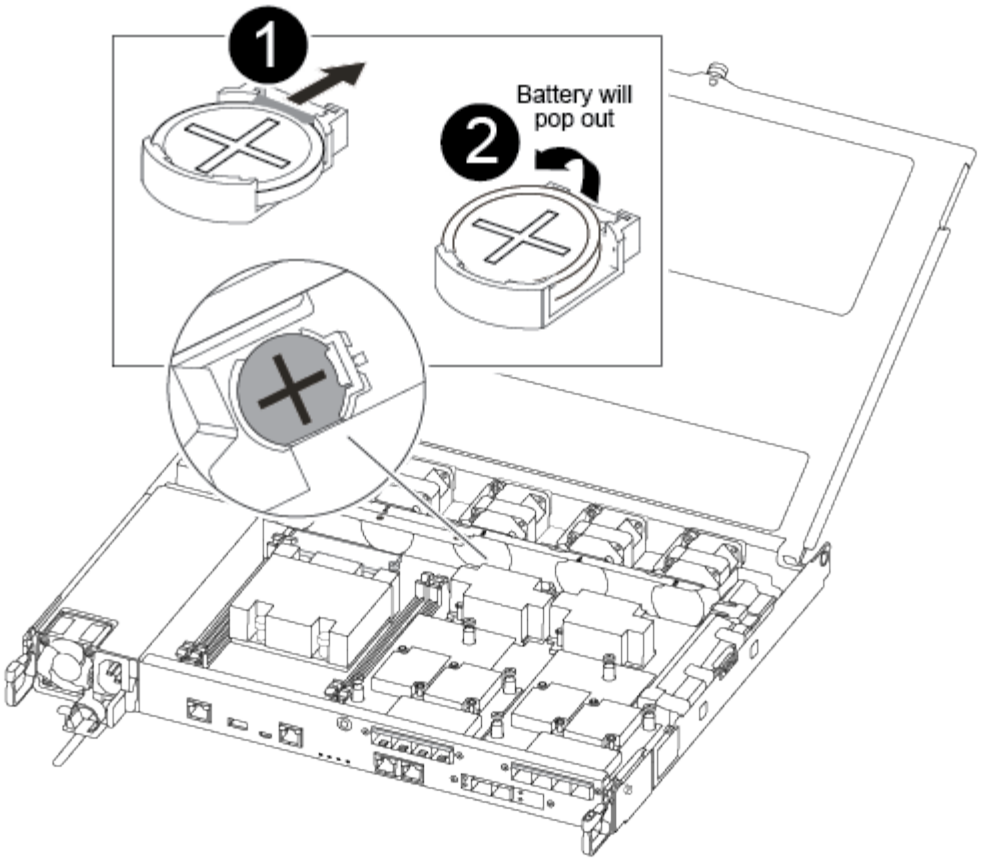

1	Gently pull tab away from the battery housing. **Attention:** Pulling it away aggressively might displace the tab.
2	Lift the battery up. **Note:** Make a note of the polarity of the battery.
3	The battery should eject out.

2. Remove the replacement battery from the antistatic shipping bag.
3. Locate the RTC battery holder between the heatsink and the midplane and insert it exactly as shown in the graphic.

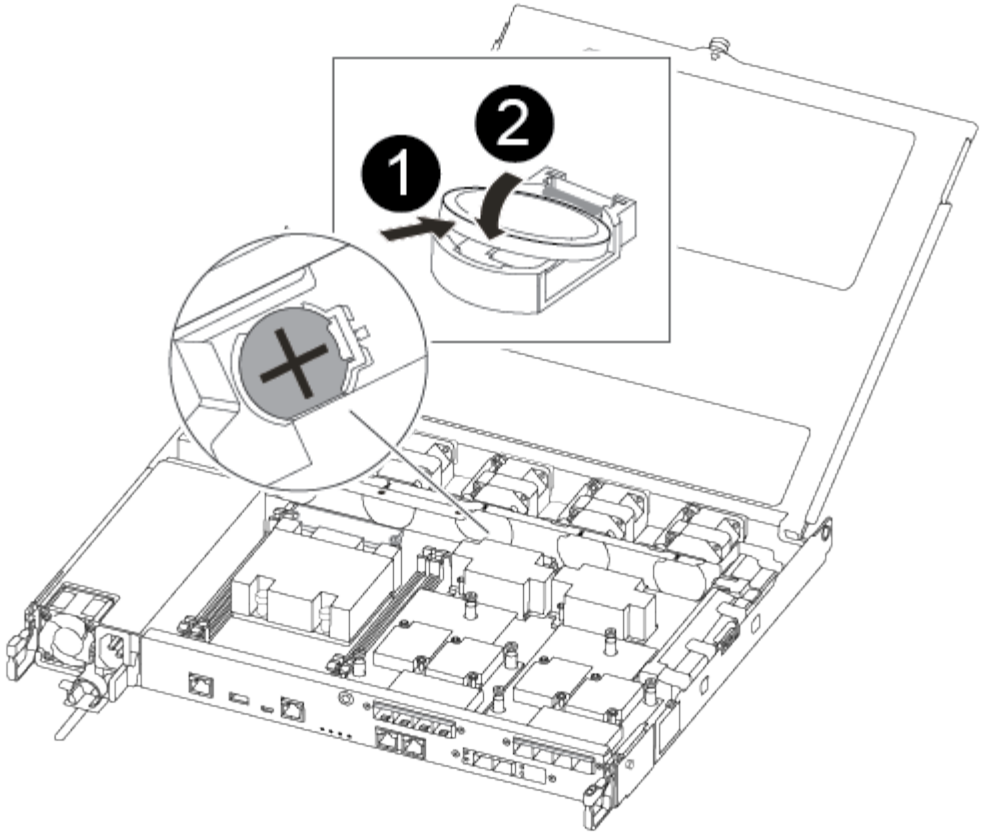

1	With positive polarity face up, slide the battery under the tab of the battery housing.
2	Push the battery gently into place and make sure the tab secures it to the housing. Pushing it in aggressively might cause the battery to eject out again.

4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Insert the controller module into the chassis:
    a. Ensure the latching mechanism arms are locked in the fully extended position.
    b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
    c. Place your index fingers through the finger holes from the inside of the latching mechanism.
    d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
    e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

       The controller module should be fully inserted and flush with the edges of the chassis.
    f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

       The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.
    g. Halt the controller at the LOADER prompt.
5. Reset the time and date on the controller:
    a. Check the date and time on the healthy controller with the `show date` command.
    b. At the LOADER prompt on the target controller, check the time and date.
    c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
    d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
    e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF A400 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

For MetroCluster configurations, see either:

- Install MetroCluster IP configuration
- Install MetroCluster Fabric-Attached configuration

Quick guide - AFF A400

The Installation and Setup instructions give graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Use the links: AFF A400 Installation and Setup Instructions.

The ASA A400 uses the same installation procedure as the AFF A400 system.

Video steps - AFF A400

The following video shows how to install and cable your new system.

Animation - AFF A400 Installation and setup instructions

Detailed guide - AFF A400

This guide gives detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your system, you need to create an account, register the system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
100 GbE cable (QSF(28)	X66211A-05 (112-00595), 0.5m X66211A-1 (112-00573), 1m X66211A-2 (112-00574), 2m X66211A-5 (112-00574), 5m		Storage, cluster interconnect/HA, and Ethernet data (order-dependent)
25 GbE cable (SFP28s)	X66240-2 (112-00598), 2m X66240-5 (112-00639), 5m		GbE network connection (order-dependent)
32 Gb FC (SFP+ Op)	X66250-2 (112-00342), 2m X66250-5 (112-00344), 5m X66250-15 (112-00346), 15m		FC network connection
Storage Cables	X66030A (112-00435), .5m X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		mini-SAS HD to mini-SAS HD cables (order-dependent)
Optical cables	X66250-2-N-C (112-00342)		16 Gb FC or 25GbE cables for mezzanine cards (order-dependent)

Type of cable…	Part number and length	Connector type	For…
RJ-45 (order dependent)	X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network
Micro-USB console cable	Not applicable		Console connection used during software setup if laptop or console does not support network discovery.
Power cables	Not applicable		Powering up the system

4. Review the *NetApp ONTAP Configuration Guide* and collect the required information listed in that guide.

   ONTAP Configuration Guide

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

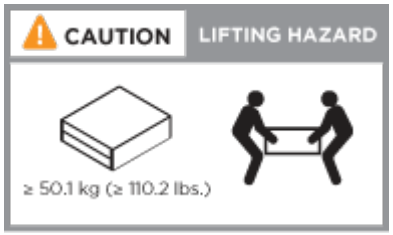

3. Attach cable management devices (as shown).

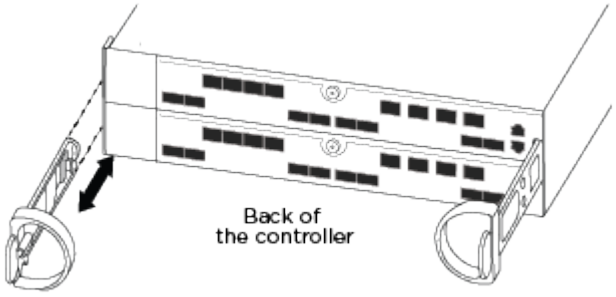

4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

If the port labels on the card are not visible, check the card installation orientation (the PCIe connector socket is on the left side of the card slot in the A400 and FAS8300/8700), and then look for the card, by part number, in the NetApp Hardware Universe for a graphic of the bezel which will show the port labels. The card part number can be found using the `sysconfig -a` command or on the system packing list.

If you are cabling an MetroCluster IP configuration, ports e0a/e0b are available for hosting data LIFs (usually in Default IPSpace).

Option 1: Cable a two-node switchless cluster

The optional data ports, optional NIC cards, and management ports on the controller modules are connected to switches. The cluster interconnect and HA ports are cabled on both controller modules.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.

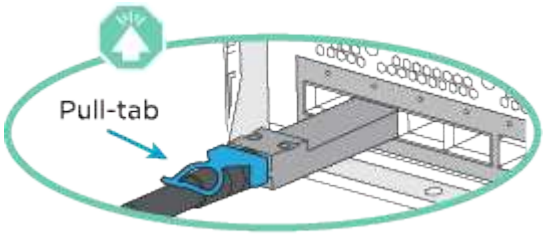

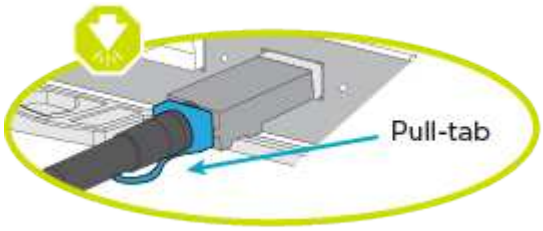

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Two-node switchless cluster cabling

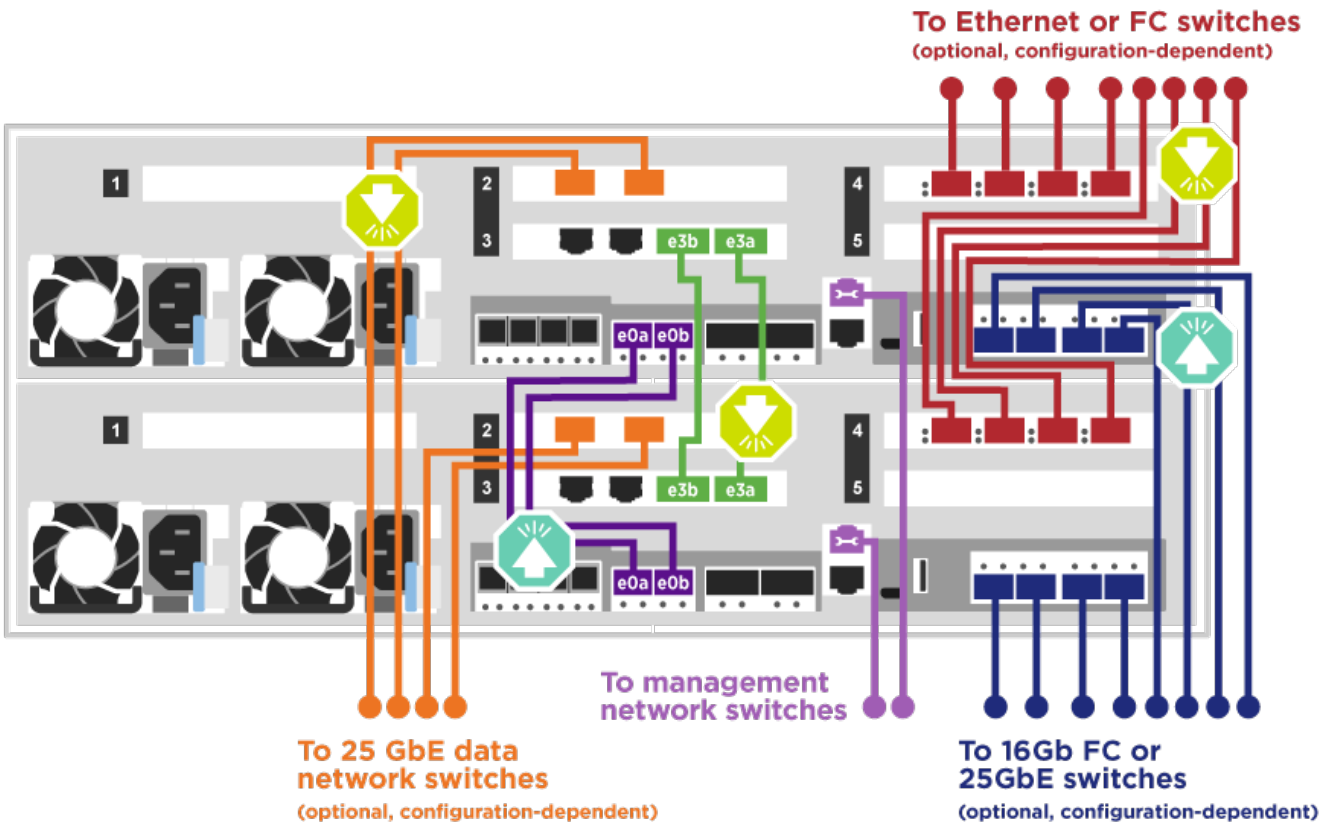

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Option 2: Cable a switched cluster

The optional data ports, optional NIC cards, mezzanine cards, and management ports on the controller modules are connected to switches. The cluster interconnect and HA ports are cabled on to the cluster/HA switch.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.

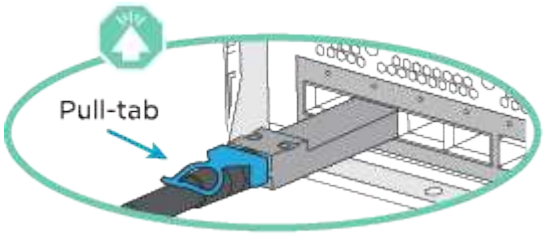

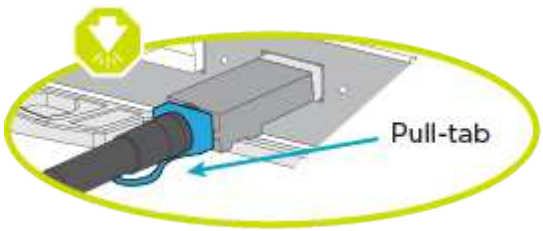

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Switched cluster cabling

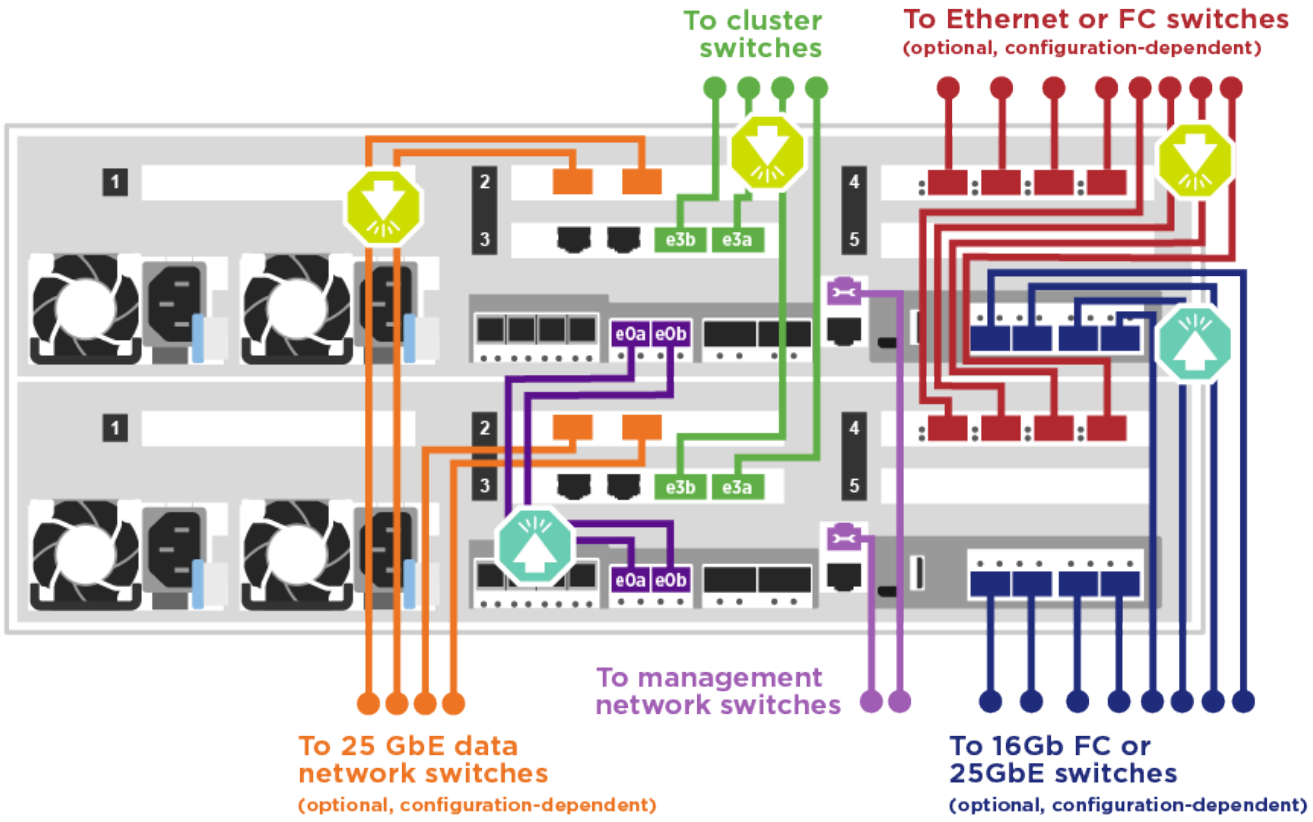

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Step 4: Cable controllers to drive shelves

You can cable either NSS224 or SAS shelves to you system.

Option 1: Cable the controllers to a single drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the NS224 are up.

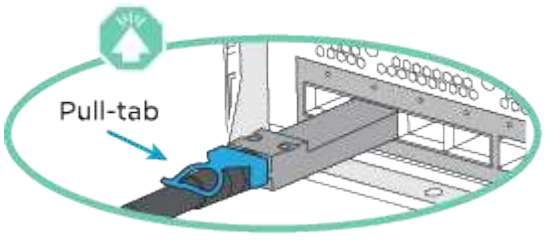

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following animation or illustration to cable your controllers to a single drive shelf.

   Animation - Cable the controllers to one NS224 drive shelf

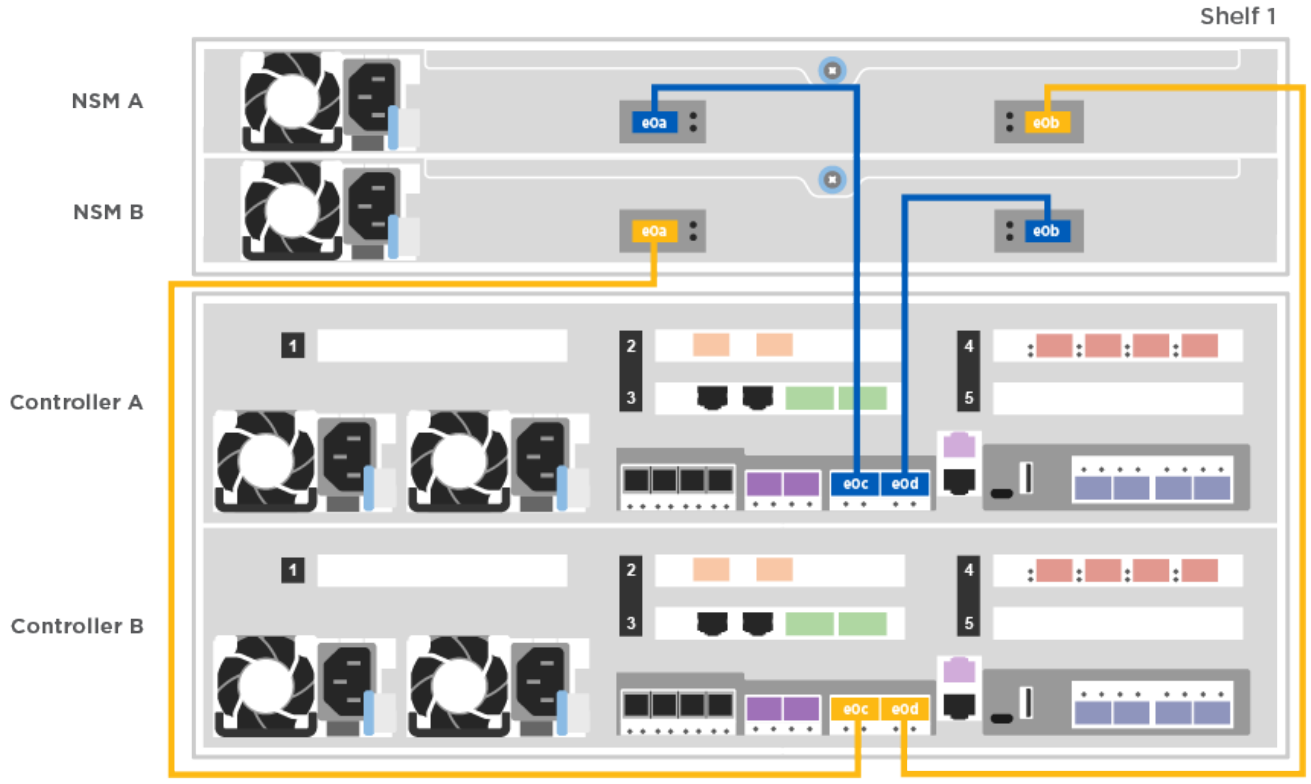

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Option 2: Cable the controllers to two drive shelves

You must cable each controller to the NSM modules on both NS224 drive shelves.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the NS224 are up.

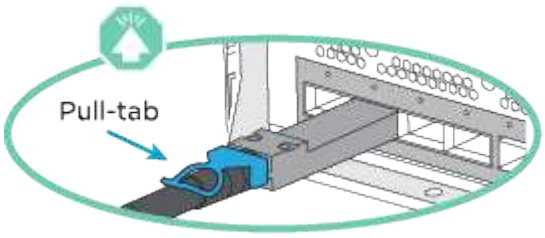

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following animation or illustration to cable your controllers to two drive shelves.

   Animation - Cable the controllers to one NS224 drive shelf

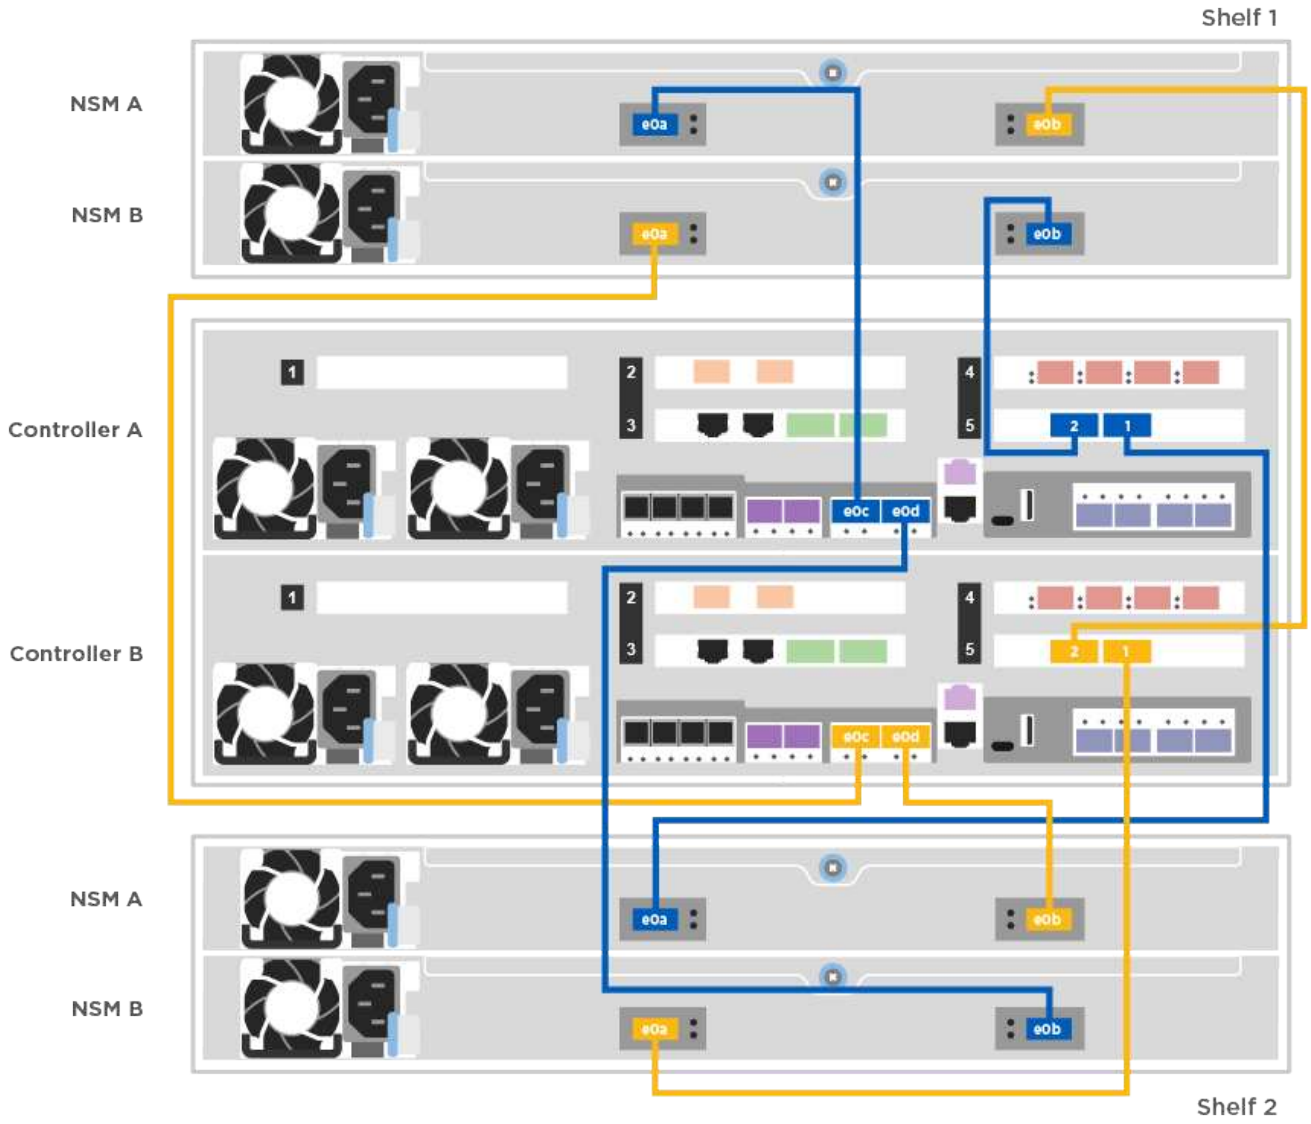

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Option 3: Cable the controllers to SAS drive shelves

You must cable each controller to the IOM modules on both SAS drive shelves.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the DS224-C are down.

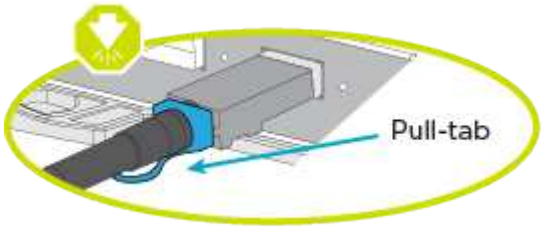

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following illustration to cable your controllers to two drive shelves.

   Animation - Cable the controllers to SAS drive shelves

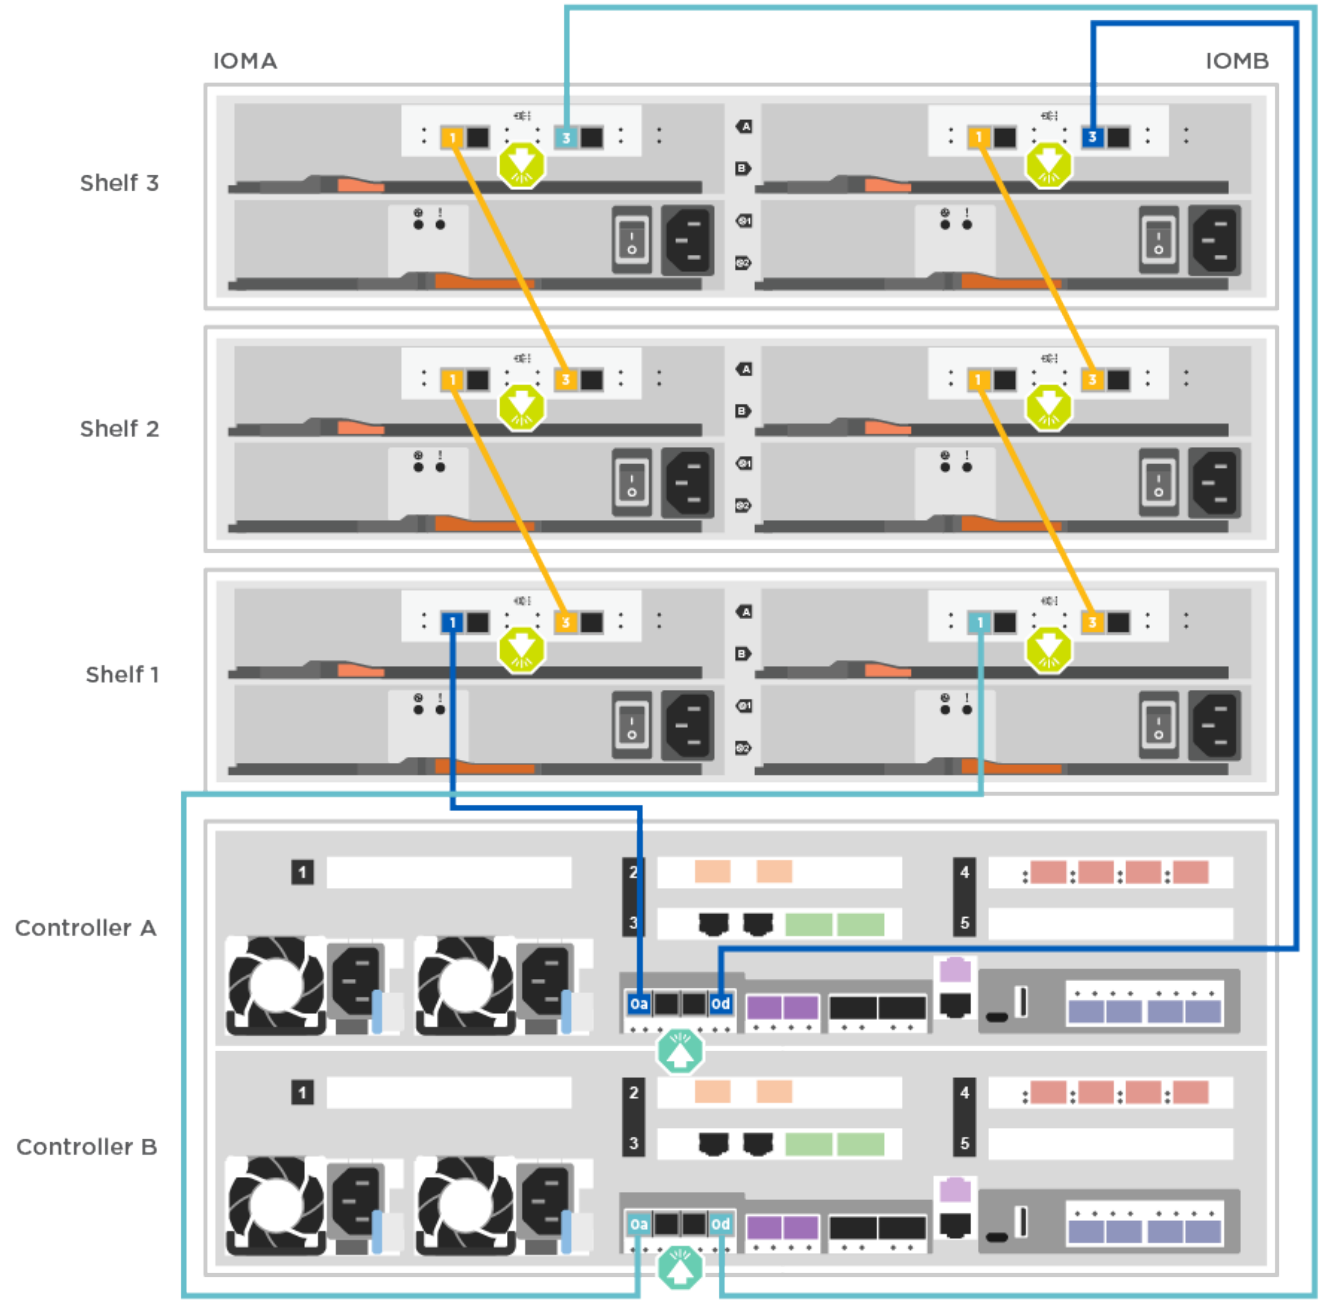

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Completing system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

1. Use the following animation to set one or more drive shelf IDs:

   If your system has NS224 drive shelves, the shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located.

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

5. Select an ONTAP icon listed to discover:

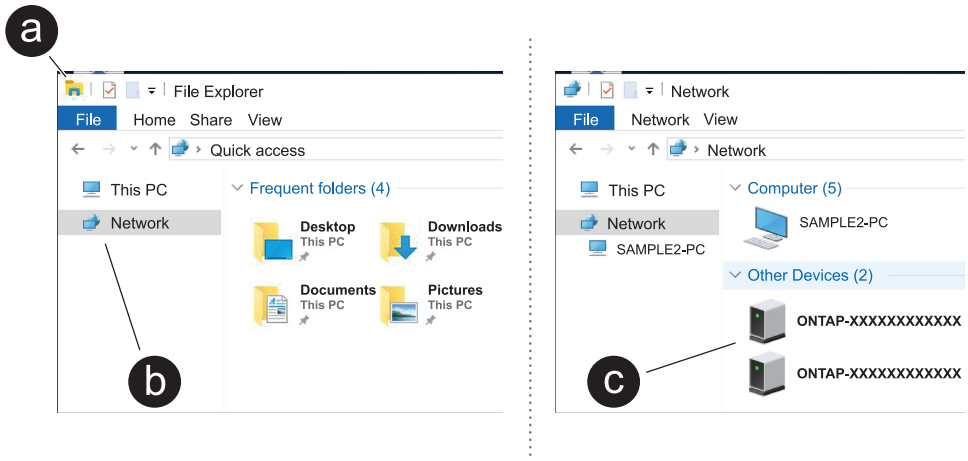

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

   XXXXX is the system serial number for the target node.

   System Manager opens.

6. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

   ONTAP Configuration Guide

7. Set up your account and download Active IQ Config Advisor:

a. Log in to your existing account or create an account.

   NetApp Support Registration

b. Register your system.

   NetApp Product Registration

c. Download Active IQ Config Advisor.

   NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet .
   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Use the following animation to set one or more drive shelf IDs:

   Animation - Set drive shelf IDs

   If your system has NS224 drive shelves, the shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located.

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   FAS8300 and FAS8700 shown.

   Animation - Power on the controllers

   Initial booting may take up to eight minutes.

4. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.
      The format for the address is https://x.x.x.x.
   b. Configure the system using the data you collected in the *NetApp ONTAP Configuration guide*.
      ONTAP Configuration Guide
6. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.
      NetApp Support Registration
   b. Register your system.
      NetApp Product Registration
   c. Download Active IQ Config Advisor.
      NetApp Downloads: Config Advisor
7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain AFF A400 hardware

Maintain the hardware of your AFF A400 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the AFF A400 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF A400 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.
NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.
PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - AFF A400

Get started with replacing the boot media in your AFF A400 storage system by reviewing the replacement requirements, shutting down the controller, replacing the boot media, restoring the image on the boot media, and verifying the system functionality.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - AFF A400

Before replacing the boot media in your AFF A400, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.

- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - AFF A400

Shut down the impaired controller in your AFF A400 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - AFF A400

The boot media in your AFF A400 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

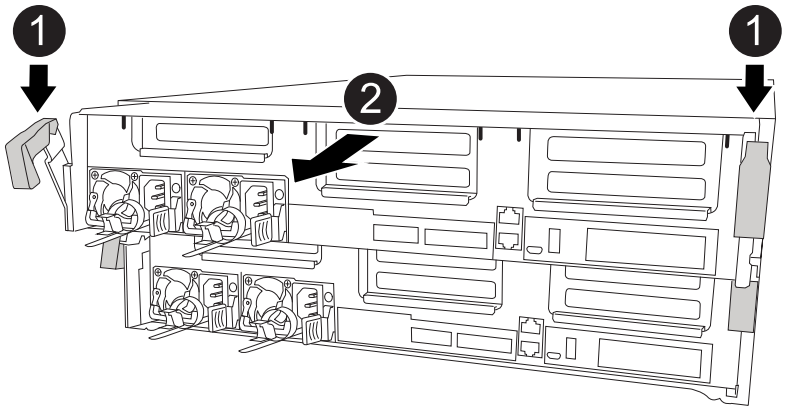

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. Open the air duct:

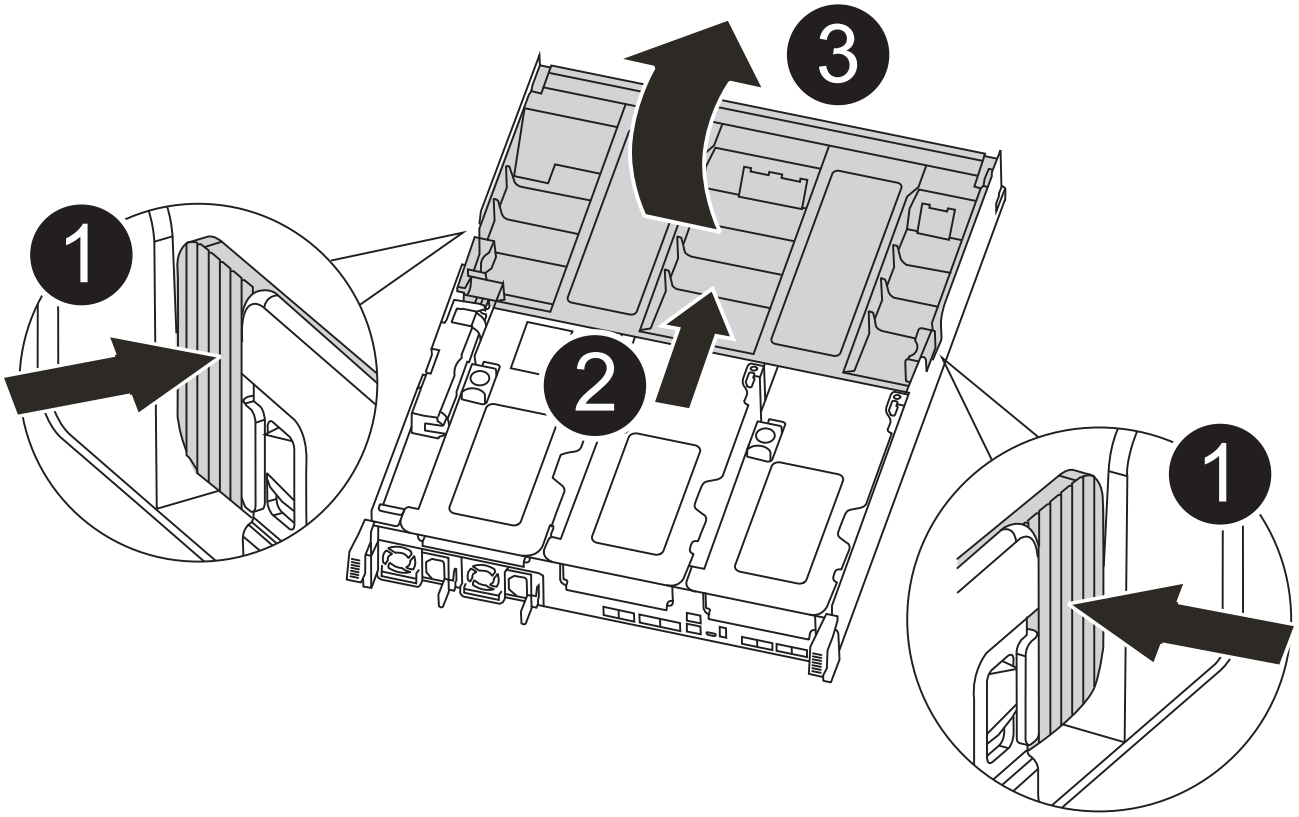

1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.

b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

9. Locate and remove the boot media from the controller module:

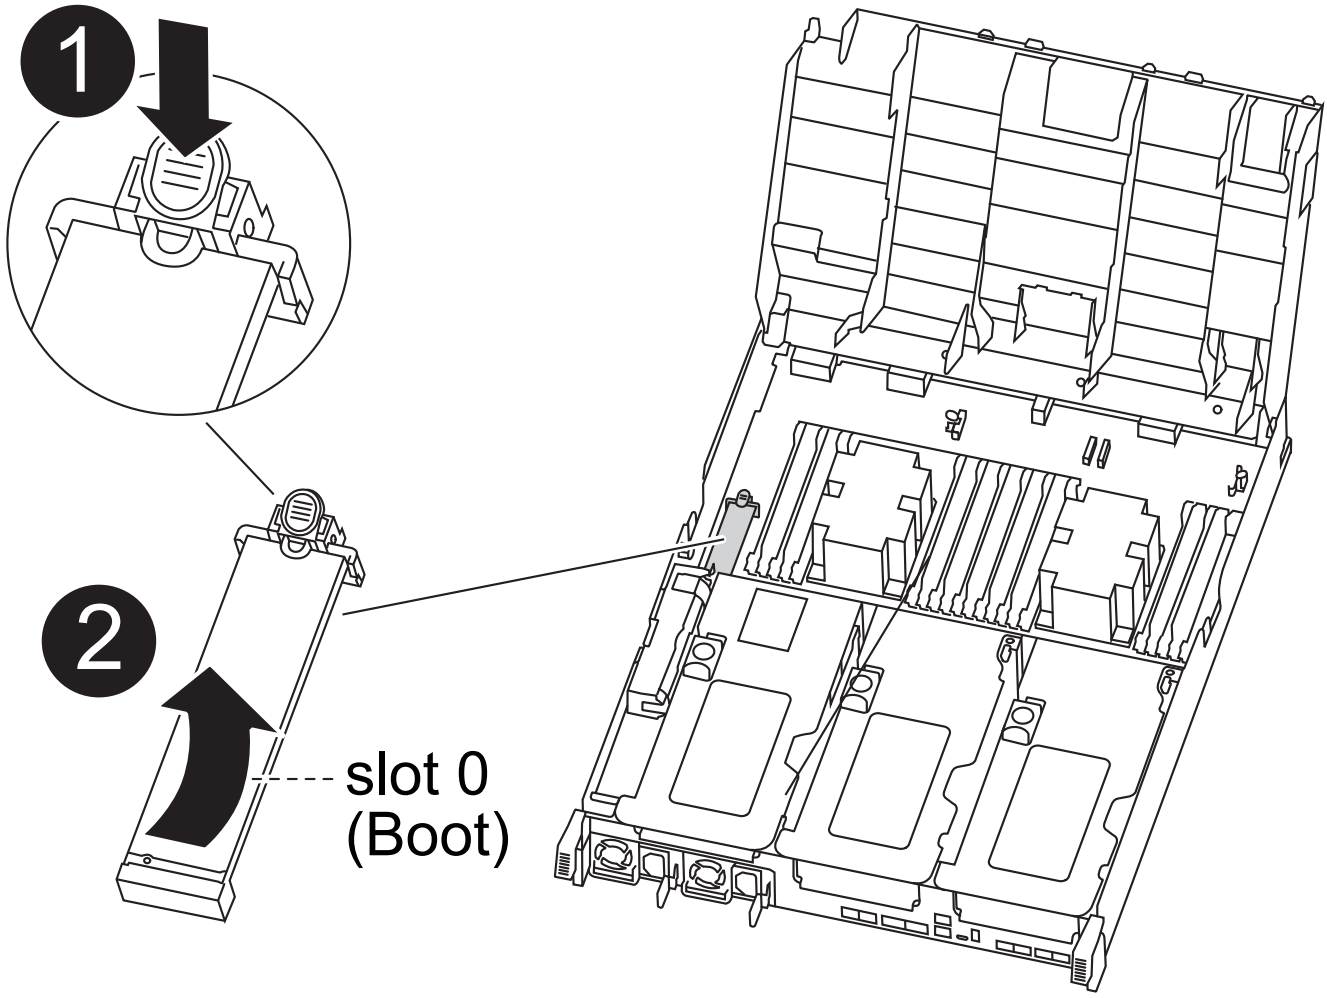

1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.

b. Rotate the boot media up and gently pull the boot media out of the socket.

10. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
11. Check the boot media to make sure that it is seated squarely and completely in the socket.

    If necessary, remove the boot media and reseat it into the socket.

12. Lock the boot media in place:
    a. Rotate the boot media down toward the motherboard.
    b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
    c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
13. Close the air duct.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - AFF A400

After installing the new boot media device in your AFF A400 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - AFF A400

If a component in your AFF A400 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF A400

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on the partner node to reinstall ONTAP on the replacement boot media in your AFF A400 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - AFF A400

Before replacing the boot media in your AFF A400 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - AFF A400

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - AFF A400

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster configuration

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 3: Controller is in a two-node Metrocluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes

that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the boot media and prepare for manual boot recovery - AFF A400

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

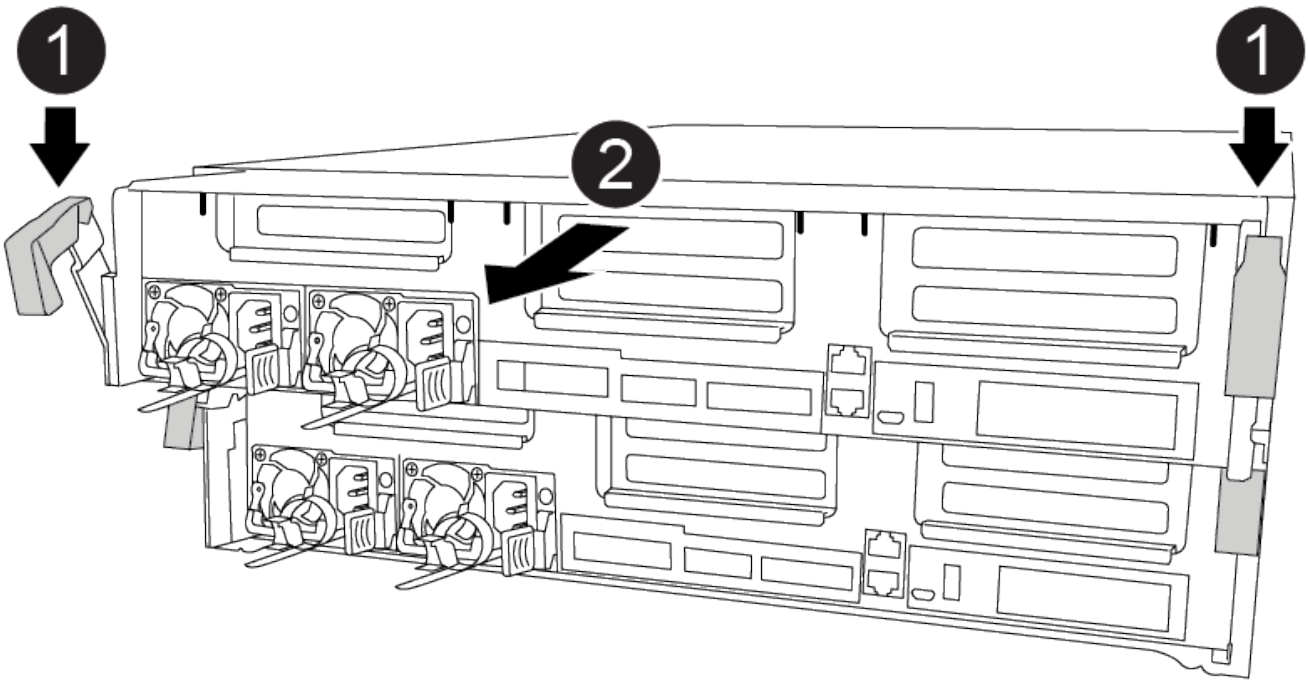

1	Locking latches

2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 2: Replace the boot media

You must locate the boot media in the controller module (see the FRU map on the controller module), and then follow the directions to replace it.

Before you begin

Although the contents of the boot media is encrypted, it is a best practice to erase the contents of the boot media before replacing it. For more information, see the Statement of Volatility for your system on the NetApp Support Site.

You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the boot media.

Animation - Replace the boot media

Steps

1. Open the air duct:

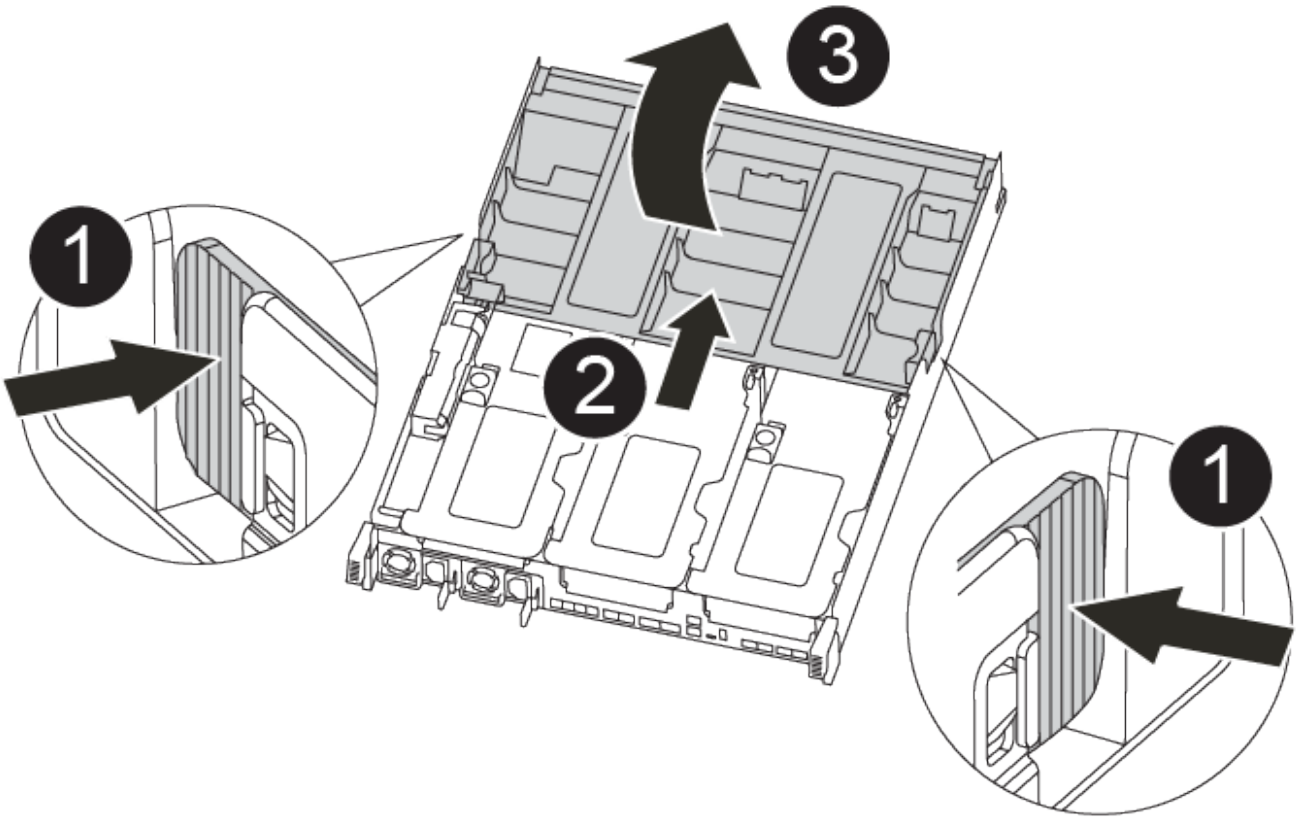

1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.

b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

2. Locate and remove the boot media from the controller module:

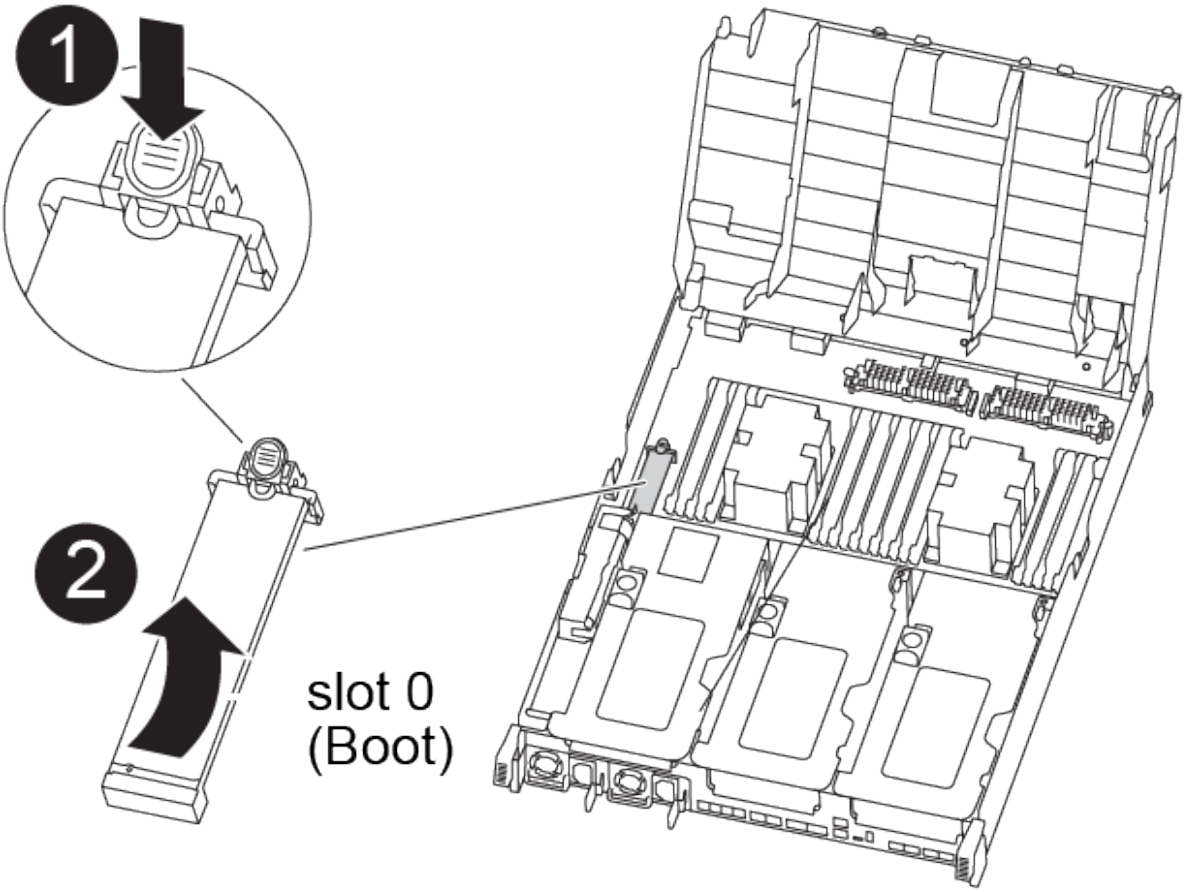

1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.

b. Rotate the boot media up and gently pull the boot media out of the socket.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the

socket.

4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Lock the boot media in place:
   a. Rotate the boot media down toward the motherboard.
   b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
   c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
6. Close the air duct.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed does not have a boot image, so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the `var` file system.

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      There are two folders in the unzipped service image file:

      - `boot`
      - `efi`

   c. Copy the `efi` folder to the top directory on the USB flash drive.

      If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

d. Remove the USB flash drive from your laptop.

2. If you have not already done so, close the air duct.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.

5. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

6. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   d. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   e. If you have not already done so, reinstall the cable management device.
7. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then `halt` the controller to boot to LOADER.

8. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t` *`initiator adapter_name`*
   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - AFF A400

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore encryption - AFF A400

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - AFF A400

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF A400

To replace the chassis, you must move the fans and controller modules from the impaired chassis to the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial

outage in a multinode cluster.

Shut down the controllers - AFF A400

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Shut down the controllers when replacing a chassis

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Shut down a controller in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the

`-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Move the fans, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move the fans

To move the fan modules to the replacement chassis when replacing the chassis, you must perform a specific sequence of tasks.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.
4. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.
5. Set the fan module aside.
6. Repeat the preceding steps for any remaining fan modules.

7. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.
10. Repeat these steps for the remaining fan modules.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - AFF A400

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for *HA-state* can be one of the following:

      - `ha`
      - `mcc`
      - `mcc-2n`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

4. Reinstall the bezel on the front of the system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - AFF A400

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement node* is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - AFF A400

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the controller module hardware - AFF A400

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following , illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

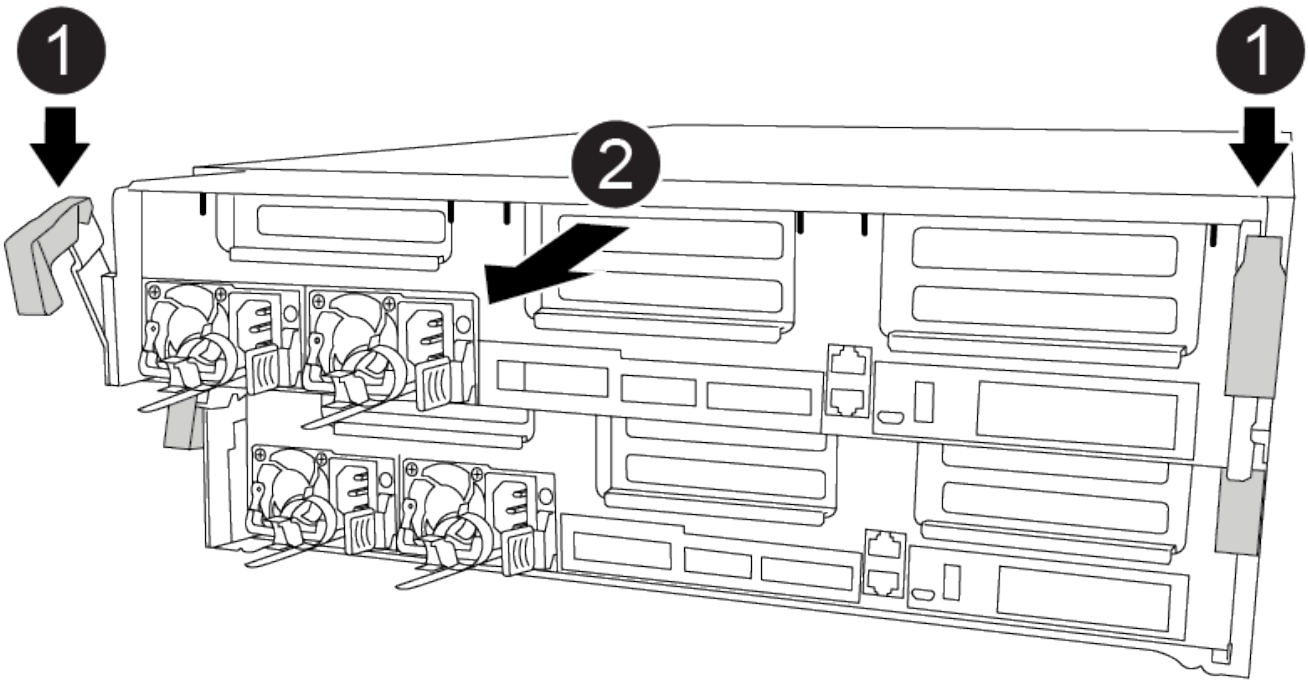

1	Locking latches

2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. On the replacement controller module, open the air duct and remove the empty risers from the controller module using the animation, illustration, or the written steps:

   Animation - Remove the empty risers from the replacement controller module

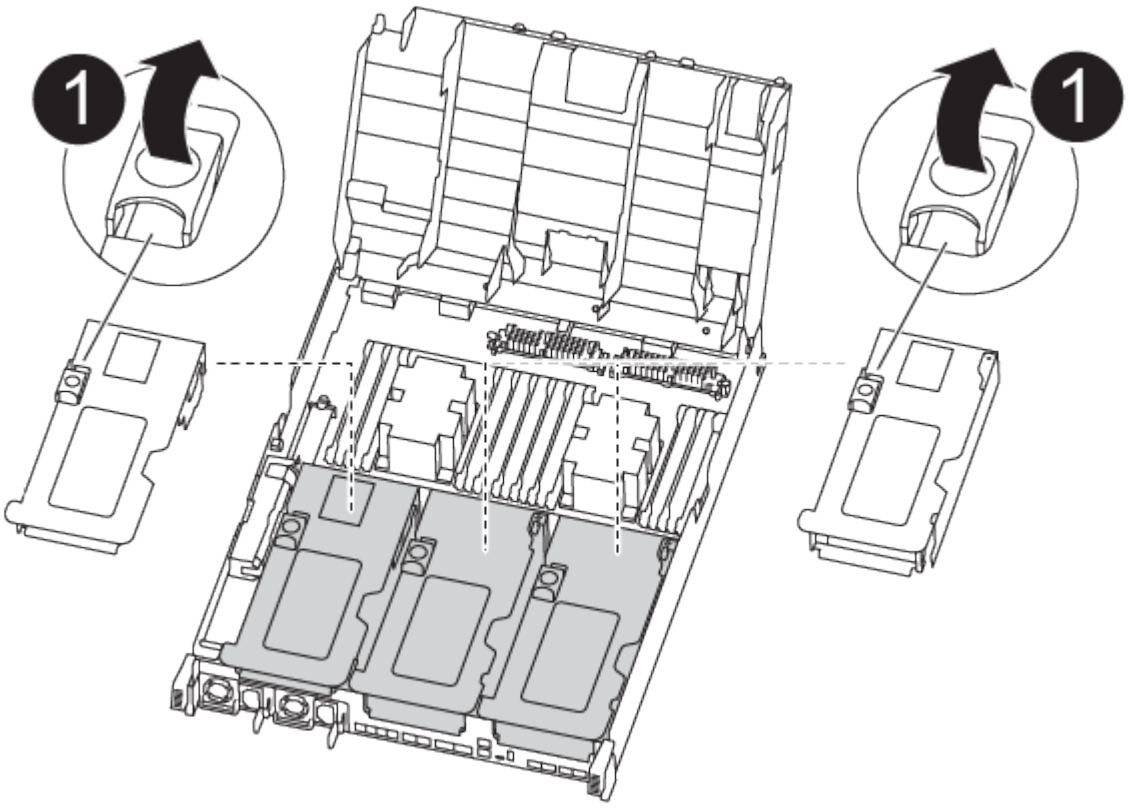

1	Riser release latches

   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
   c. Rotate the riser locking latch on the left side of riser 1 up and toward air duct, lift the riser up, and then set it aside.
   d. Repeat the previous step for the remaining risers.

Step 2: Move the power supplies

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

You can use the following animation, illustration, or the written steps to move the power supplies to the replacement controller module.

Animation - Move the power supplies

1. Remove the power supply:

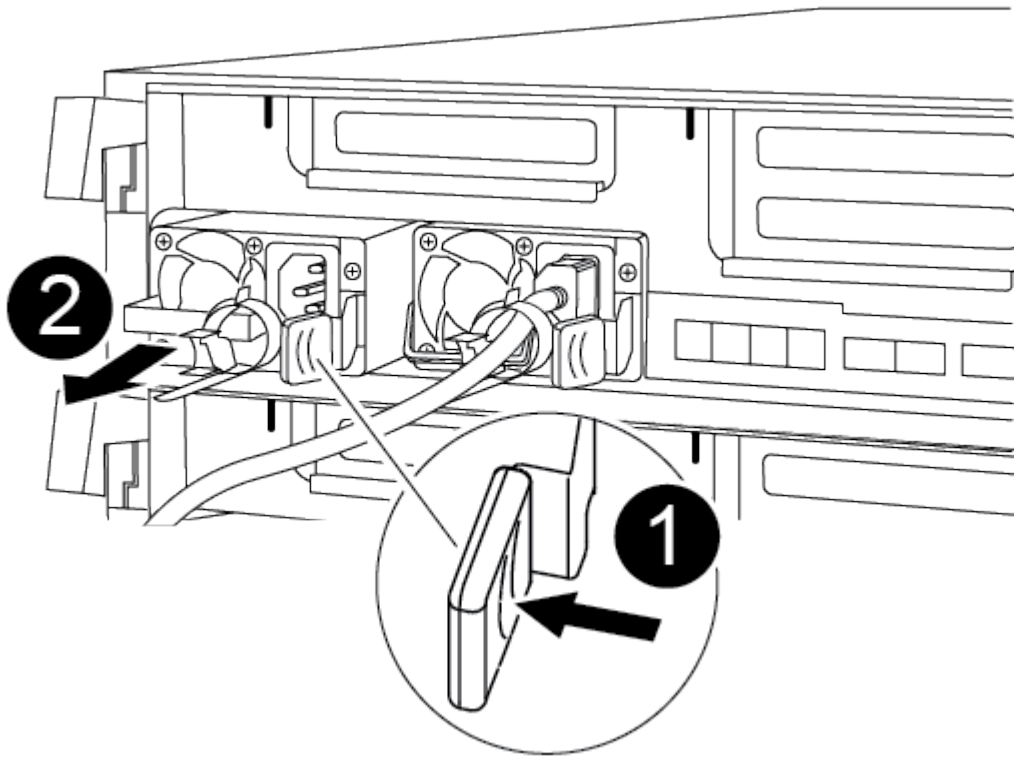

1	PSU locking tab
2	Power cable retainer

   a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
   b. Press the blue locking tab to release the power supply from the chassis.
   c. Using both hands, pull the power supply out of the chassis, and then set it aside.
2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

4. Repeat the preceding steps for any remaining power supplies.

Step 3: Move the NVDIMM battery

To move the NVDIMM battery from the impaired controller module to the replacement controller module, you must perform a specific sequence of steps.

You can use the following animation, illustration, or the written steps to move the NVDIMM battery from the impaired controller module to the replacement controller module.

Animation - Move the NVDIMM battery

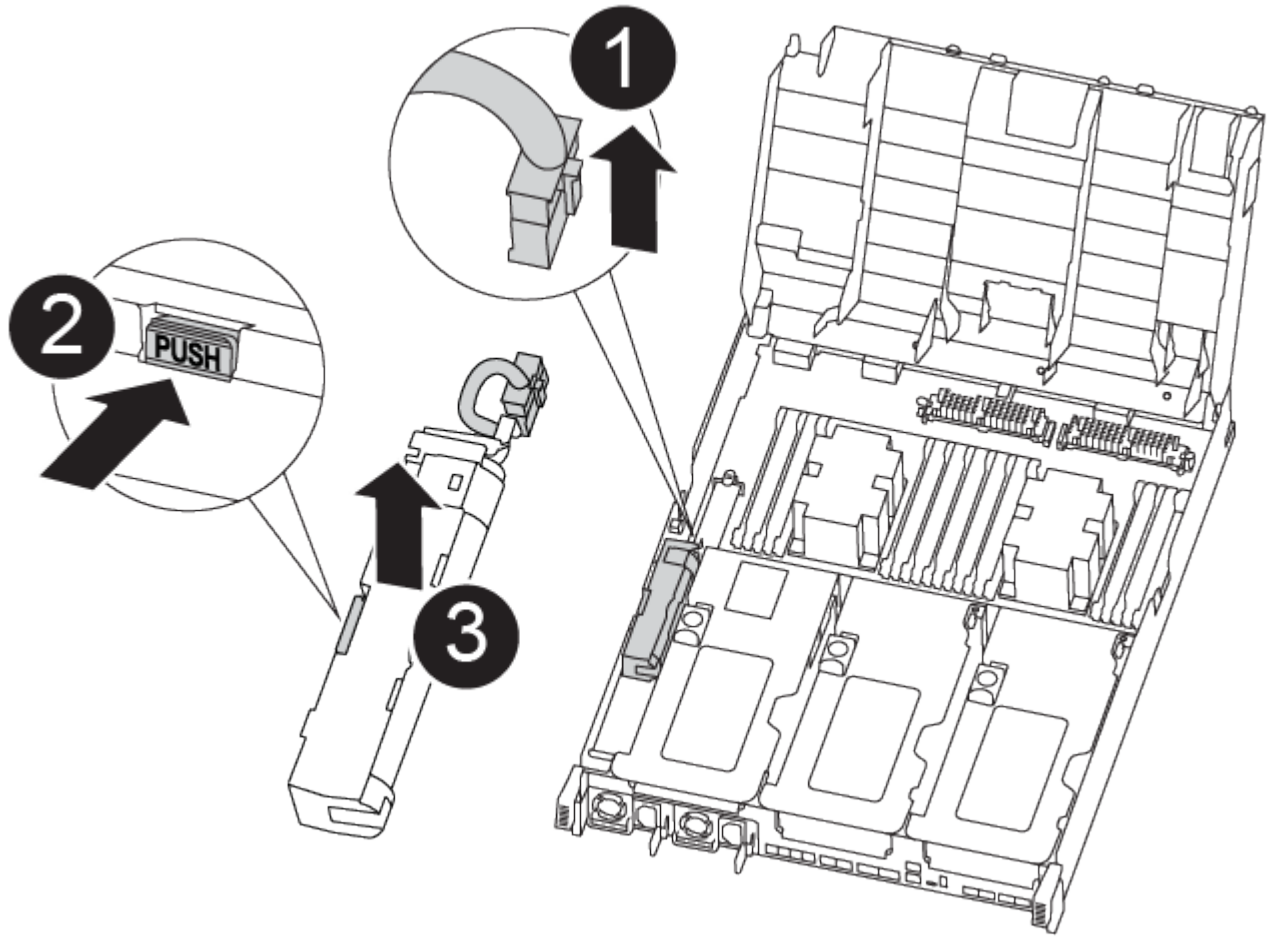

1	NVDIMM battery plug
2	NVDIMM battery locking tab
3	NVDIMM battery

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.
3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.

> Do not plug the battery cable back into the motherboard until instructed to do so.

Step 4: Move the boot media

You must locate the boot media, and then follow the directions to remove it from the impaired controller module and insert it into the replacement controller module.

You can use the following animation, illustration, or the written steps to move the boot media from the impaired controller module to the replacement controller module.

Animation - Move the boot media

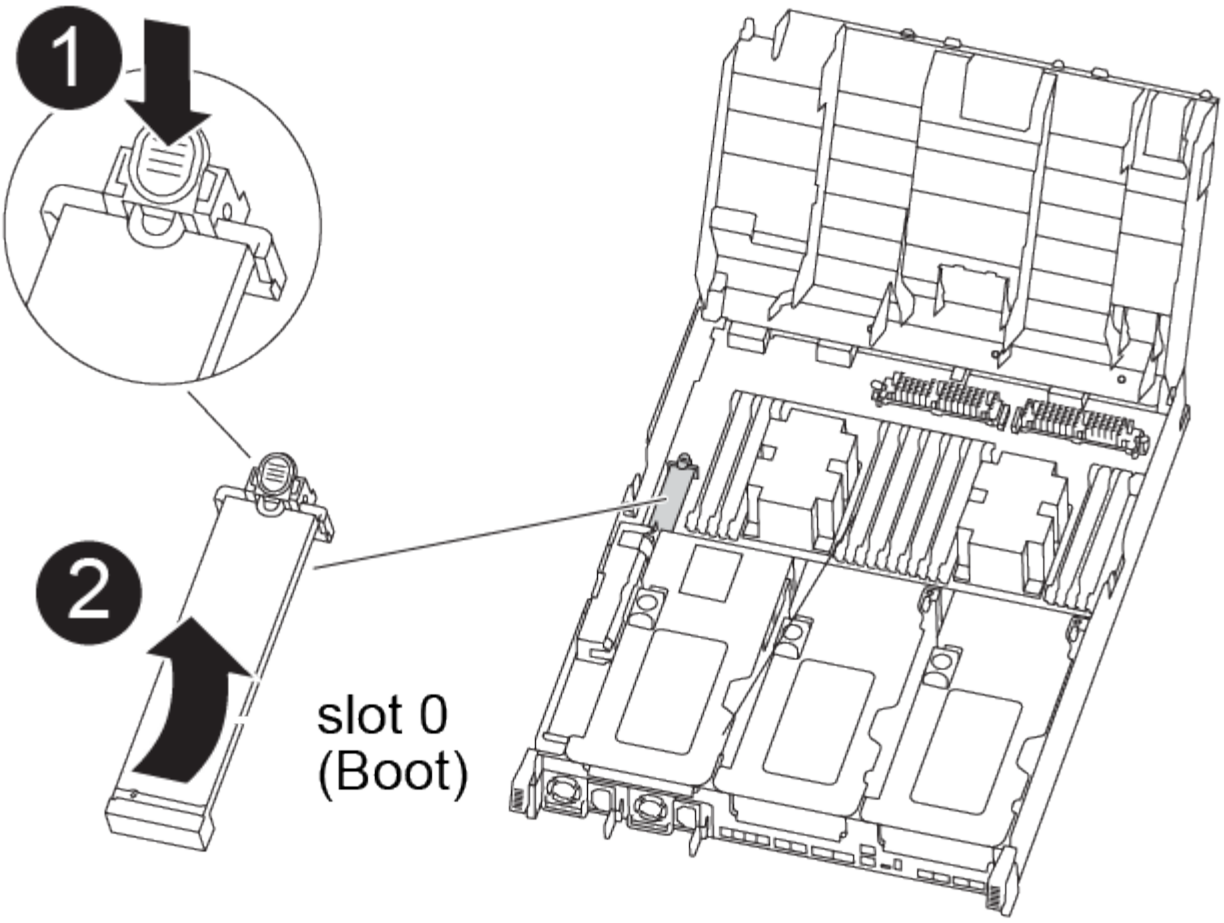

1	Boot media locking tab
2	Boot media

1. Locate and remove the boot media from the controller module:
   a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
   b. Rotate the boot media up and gently pull the boot media out of the socket.
2. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
3. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

4. Lock the boot media in place:
   a. Rotate the boot media down toward the motherboard.
   b. Press the blue locking button so that it is in the open position.
   c. Placing your fingers at the end of the boot media by the blue button, firmly push down on the boot media end to engage the blue locking button.

Step 5: Move the PCIe risers and mezzanine card

As part of the controller replacement process, you must move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

You can use the following animations, illustrations, or the written steps to move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

Moving PCIe riser 1 and 2 (left and middle risers):

Animation - Move PCI risers 1 and 2

Moving the mezzanine card and riser 3 (right riser):

Animation - Move the mezzanine card and riser 3

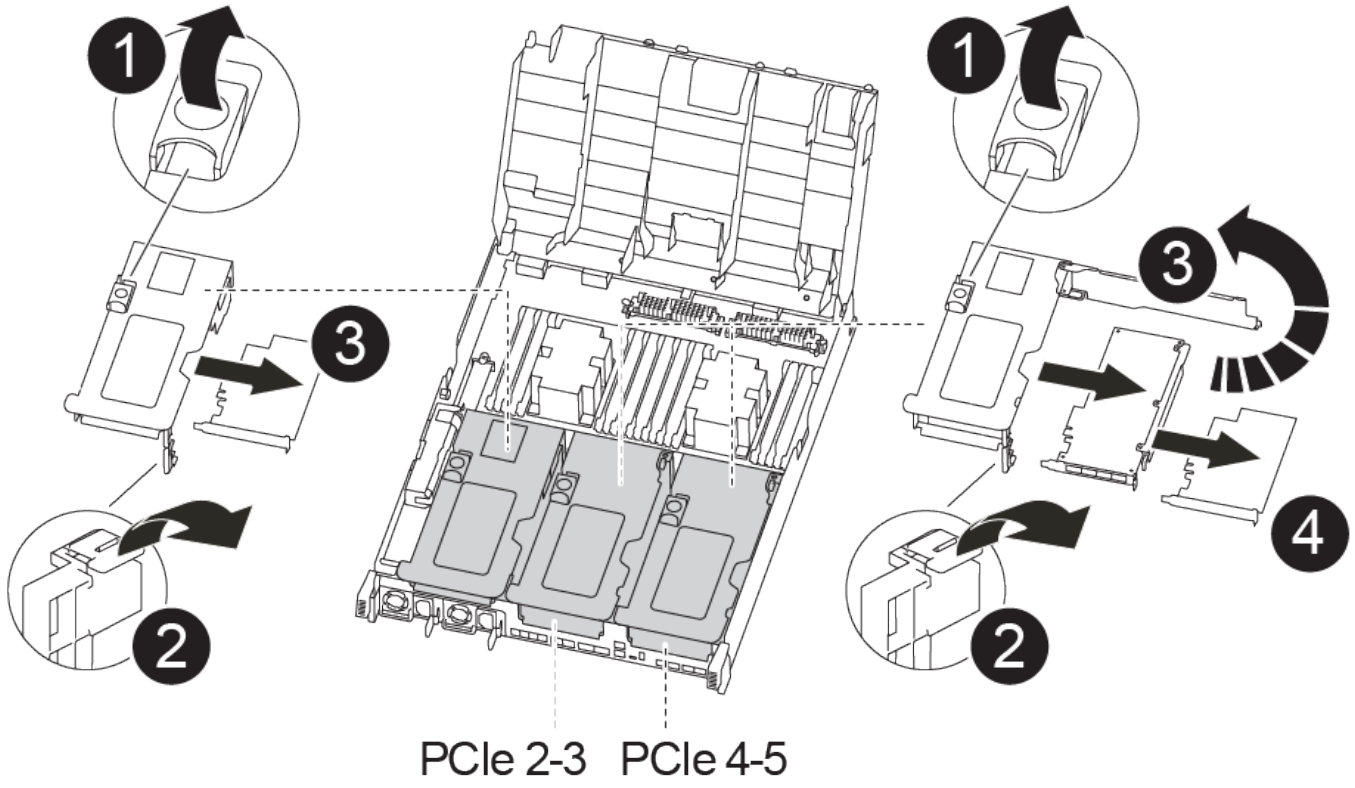

1	Riser locking latch
2	PCI card locking latch
3	PCI locking plate
4	PCI card

1. Move PCIe risers one and two from the impaired controller module to the replacement controller module:

a. Remove any SFP or QSFP modules that might be in the PCIe cards.
b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

   The riser raises up slightly from the controller module.

c. Lift the riser up, and then move it to the replacement controller module.
d. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins, push the riser squarely into the socket on the motherboard, and then rotate the latch down flush with the sheet metal on the riser.
e. Repeat this step for riser number 2.

2. Remove riser number 3, remove the mezzanine card, and install both into the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then set it aside on a stable, flat surface.
   d. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket, and then move it to the replacement controller module.
   e. Install the mezzanine in the replacement controller and secure it with the thumbscrews.
   f. Install the third riser in the replacement controller module.

Step 6: Move the DIMMs

You need to locate the DIMMs, and then move them from the impaired controller module to the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

You can use the following animation, illustration, or the written steps to move the DIMMs from the impaired controller module to the replacement controller module.

Animation - Move the DIMMs

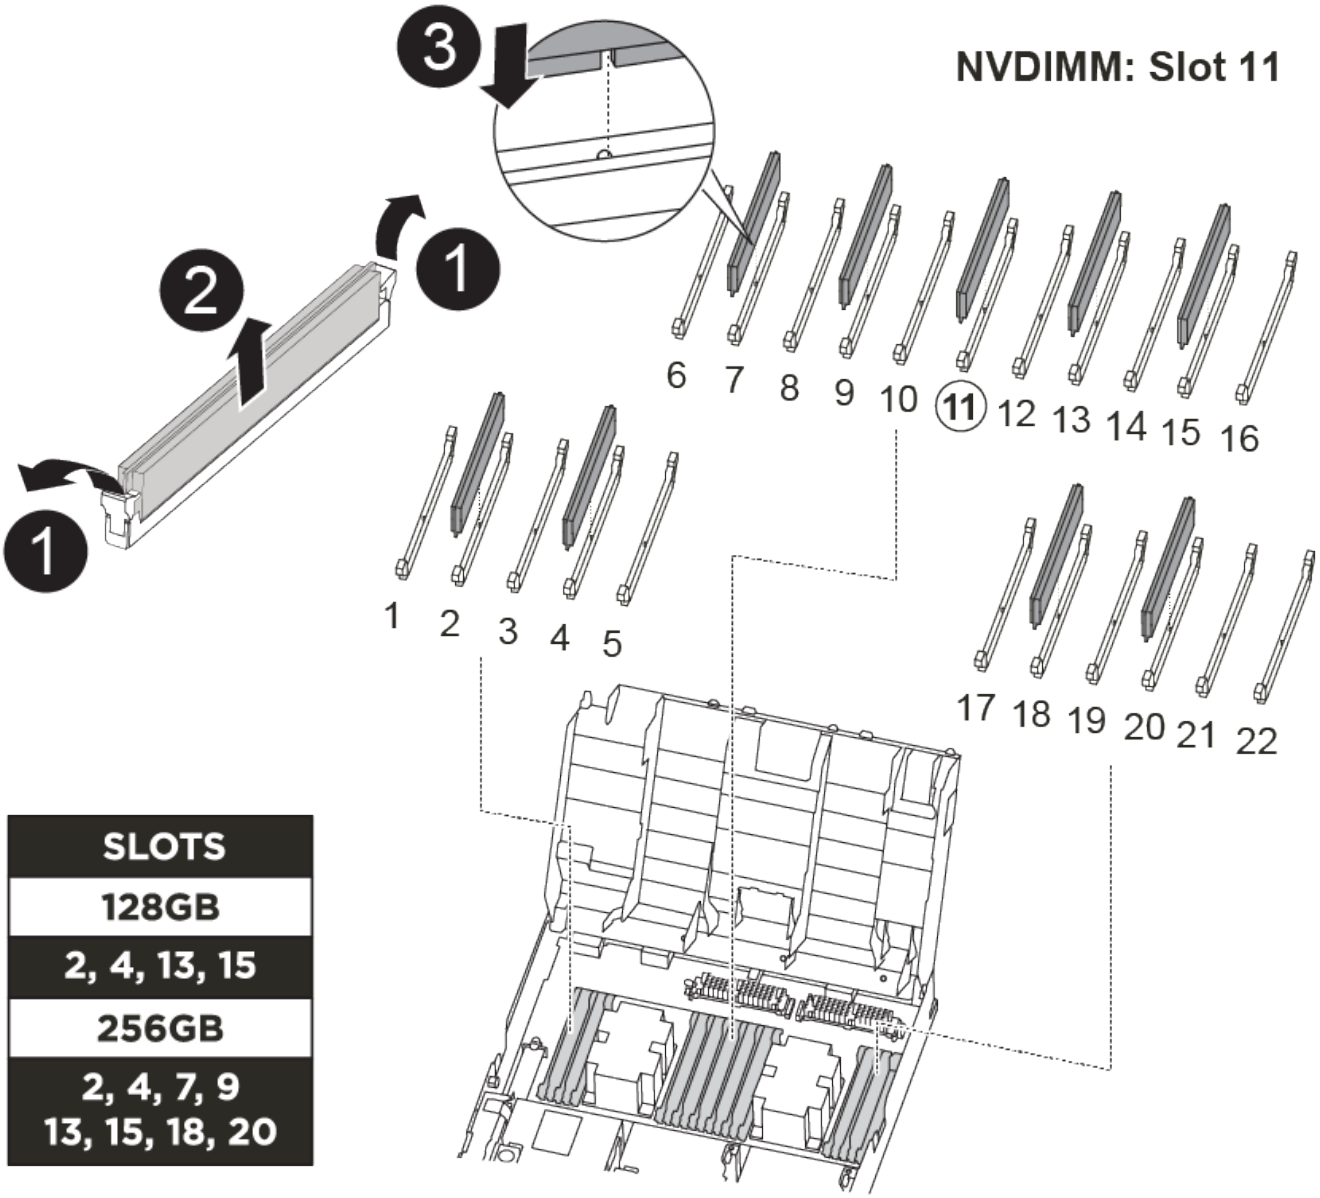

1	DIMM locking tabs
2	DIMM
3	DIMM socket

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Verify that the NVDIMM battery is not plugged into the new controller module.
4. Move the DIMMs from the impaired controller module to the replacement controller module:

Make sure that you install the each DIMM into the same slot it occupied in the impaired controller module.

a. Eject the DIMM from its slot by slowly pushing apart the DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

b. Locate the corresponding DIMM slot on the replacement controller module.

c. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the DIMM squarely into the socket.

The DIMMs fit tightly in the socket, but should go in easily. If not, realign the DIMM with the socket and reinsert it.

d. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

e. Repeat these substeps for the remaining DIMMs.

5. Plug the NVDIMM battery into the motherboard.

Make sure that the plug locks down onto the controller module.

Step 7: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

You can use the following animation, illustration, or the written steps to install the replacement controller module in the chassis.

Animation - Install the controller module

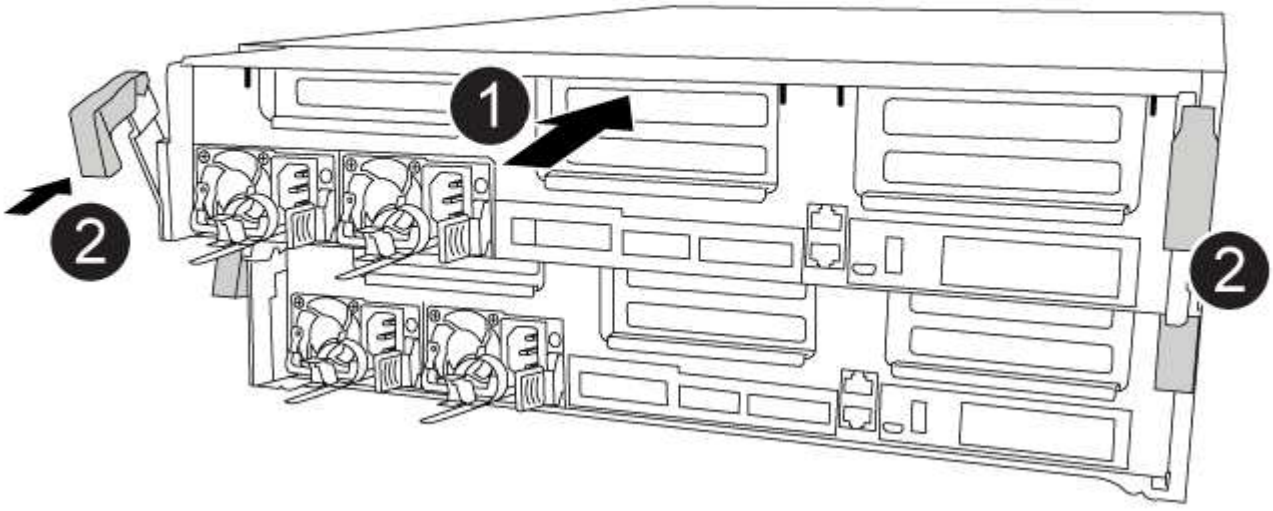

1	Controller module
2	Controller locking latches

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.
   g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

Restore and verify the system configuration - AFF A400

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF A400

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

You must use the correct procedure for your configuration:

Controller redundancy	Then use this procedure…
HA pair	Option 1: Verify the system ID change on an HA system]
Two-node MetroCluster configuration	Option 2: Manually reassign the system ID on systems in a two-node MetroCluster configuration

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and

then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Option 2: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `metrocluster node show -fields node-systemid`,dr-partner-systemid`

In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

```
dr-group-id cluster         node                 node-systemid dr-
partner-systemid
 ----------- -------------------- ------------------- -------------
-------------------
 1           Cluster_A            Node_A_1             536872914
118073209
 1           Cluster_B            Node_B_1             118073209
536872914
 2 entries were displayed.
```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

```
Local System ID: 118065481
    ...
    ...
```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                POOL   SERIAL NUMBER  HOME
-------    -------------        -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Verify that the coredumps are saved: `system node run -node local-node-name partner savecore`

      If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node local-node-name partner savecore -s command`.</info>.

   c. Return to the admin privilege level: `set -privilege admin`
7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
8. Boot the *replacement* node: `boot_ontap`
9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`
10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
12. Simulate a switchover operation:

a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

   You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`

c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - AFF A400

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`

3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A400

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

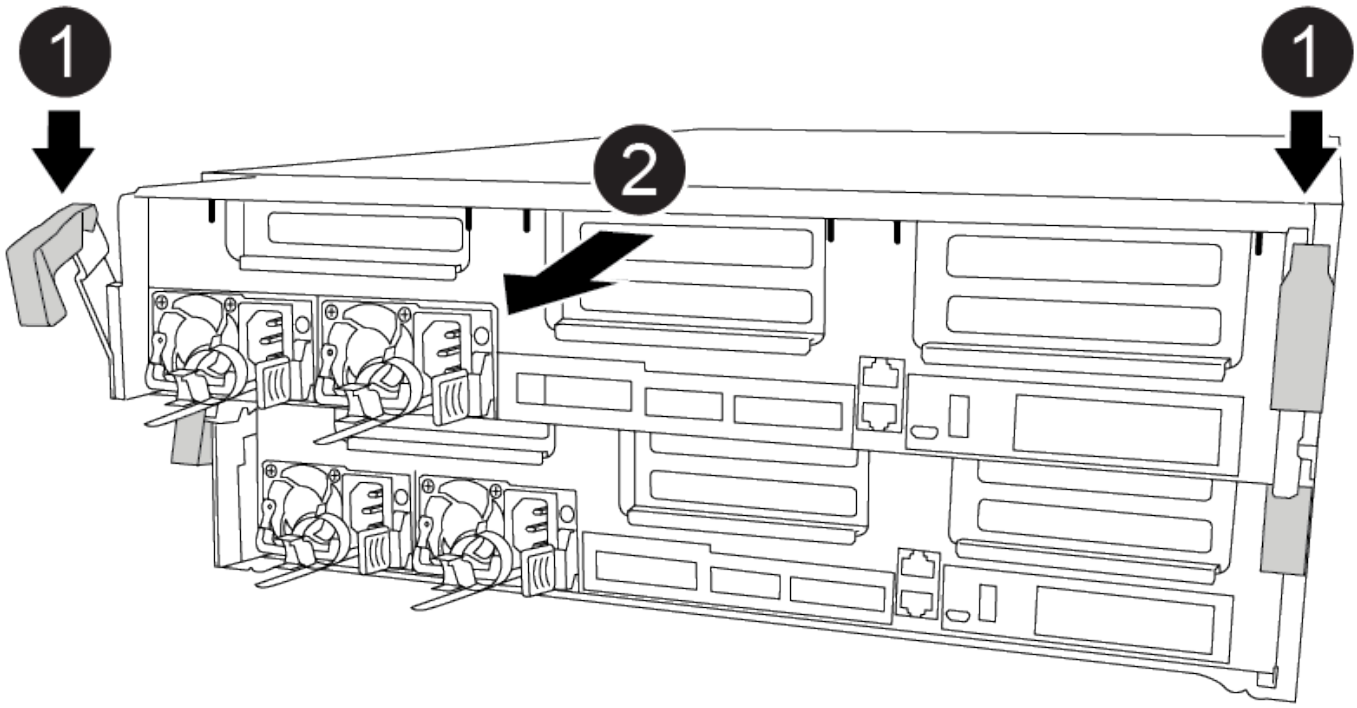

1	Locking latches
2	Controller moves slightly out of chassis

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace system DIMMs

Replacing a system DIMM involves identifying the target DIMM through the associated error message, locating the target DIMM using the FRU map on the air duct, and then replacing the DIMM.

You can use the following animation, illustration, or the written steps to replace a system DIMM.

> The animation and illustration show empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace a system DIMM

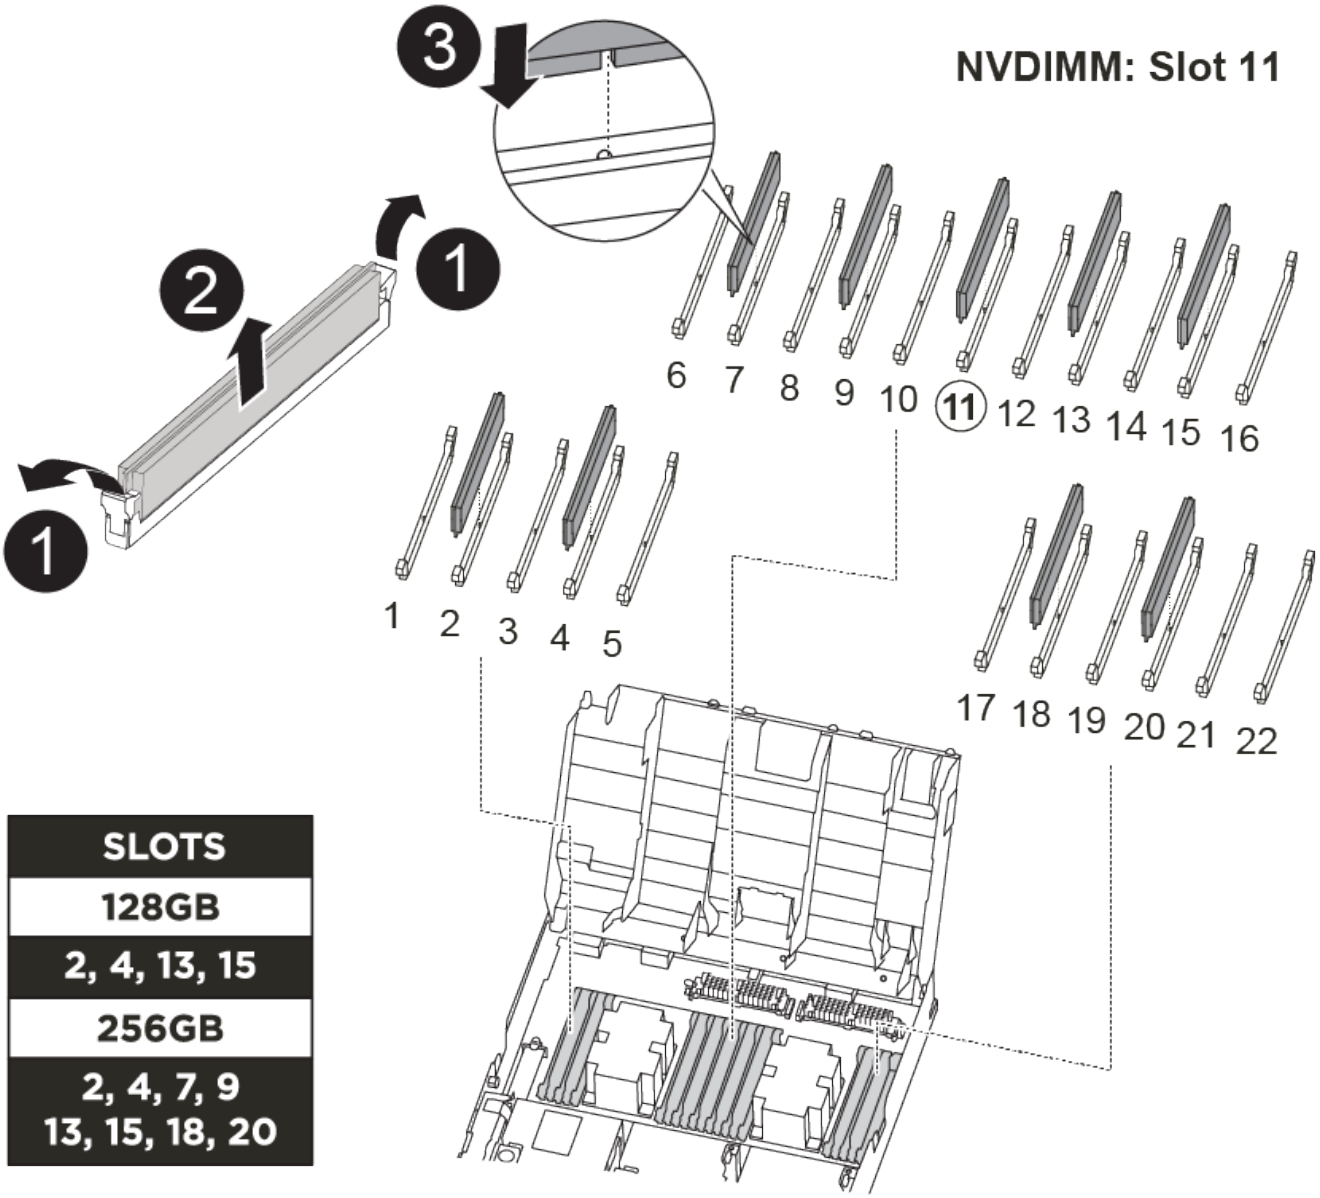

1	DIMM locking tabs
2	DIMM
3	DIMM socket

The DIMMs are located in sockets 2, 4, 13, and 15. The NVDIMM is located in slot 11.

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
4. Eject the DIMM from its socket by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the socket.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module

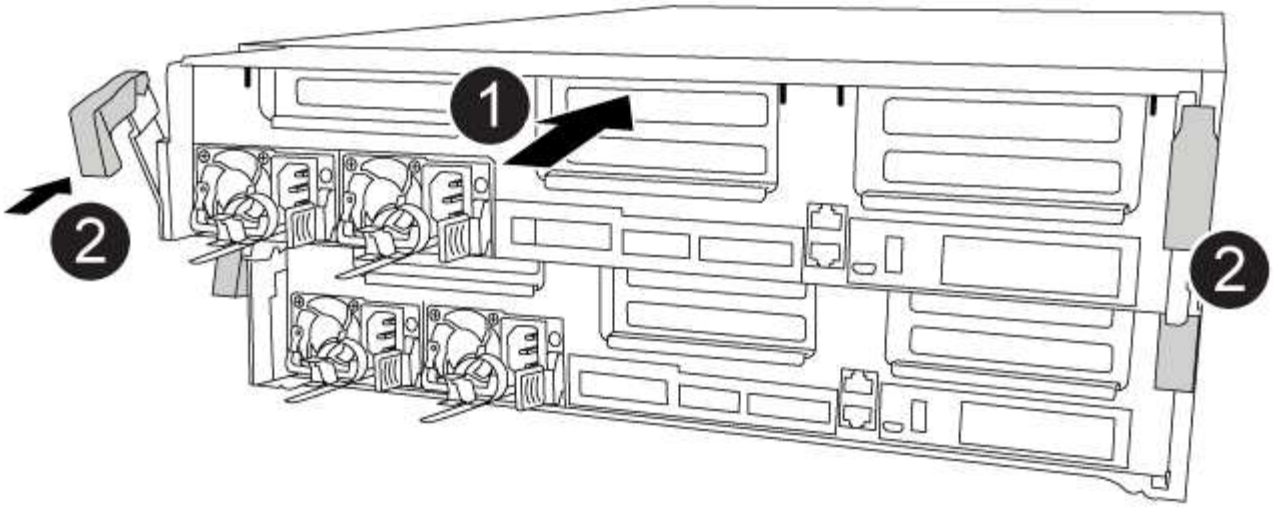

1	Controller module
2	Controller locking latches

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot

process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`

4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    -------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    -------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a fan module - AFF A400

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> ⓘ You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

You can use the following animation, illustration, or the written steps to hot-swap a fan module.

Animation - Replace a fan

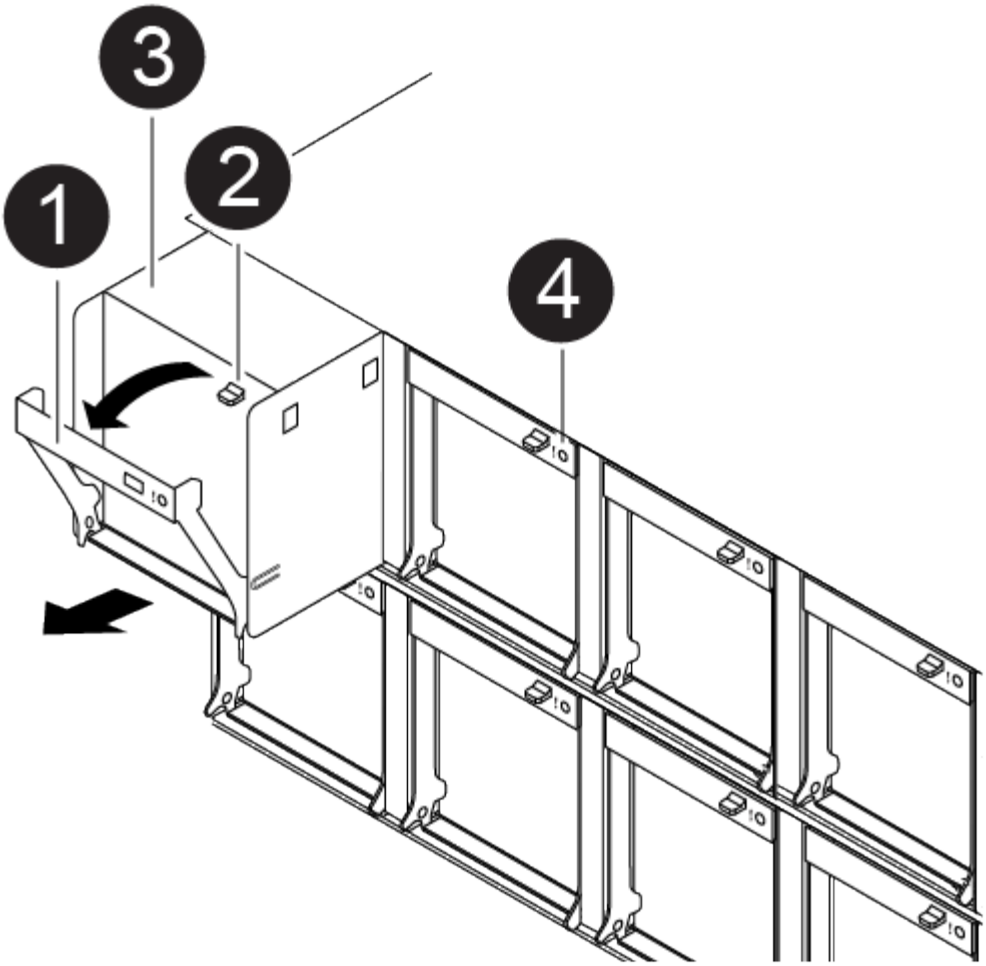

1	Fan handle
2	Locking tab
3	Fan
4	Status LED

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.

5. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

6. Set the fan module aside.
7. Insert the replacement fan module into the chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The Attention LED should not be lit after the fan is seated and has spun up to operational speed.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - AFF A400

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animations, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

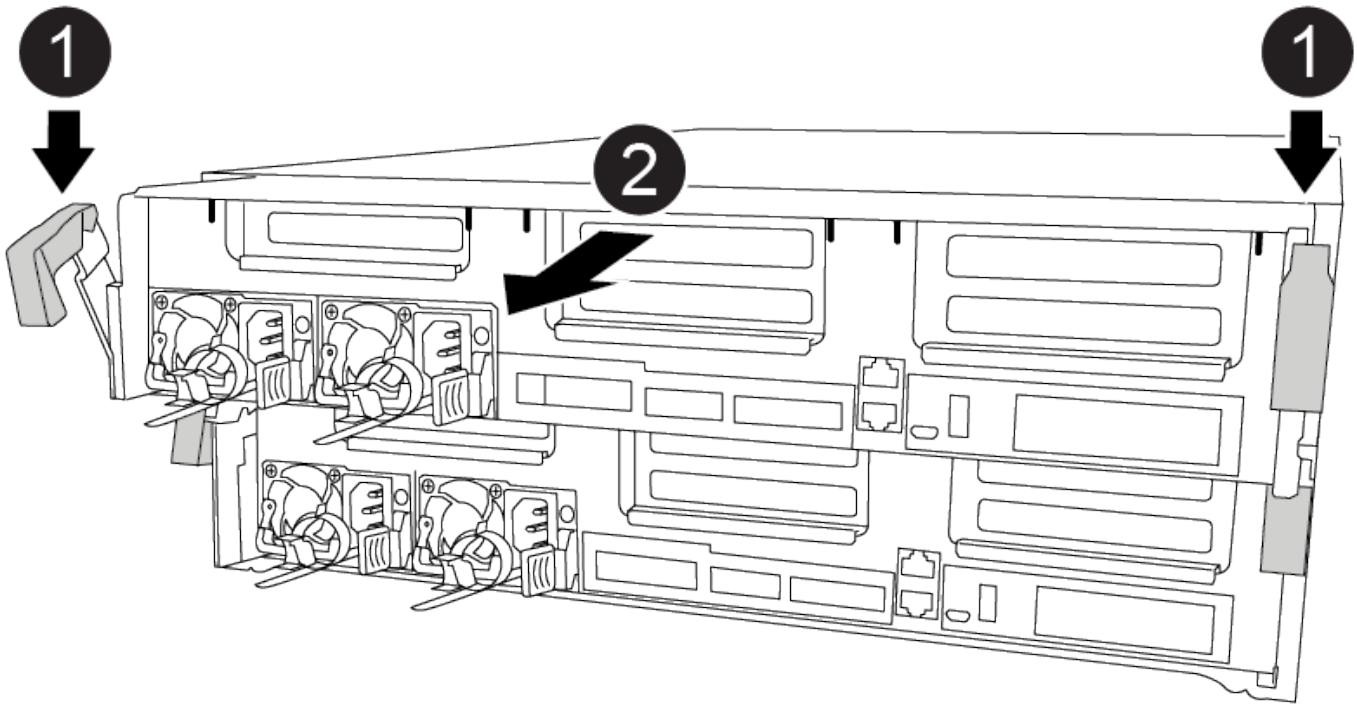

1	Locking latches
2	Controller moves slightly out of chassis

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module. See the FRU map inside the controller module to locate the NVDIMM battery.

The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.

You can use the following animation, illustration, or the written steps to replace the NVDIMM battery.

Animation - Replace the NVDIMM battery

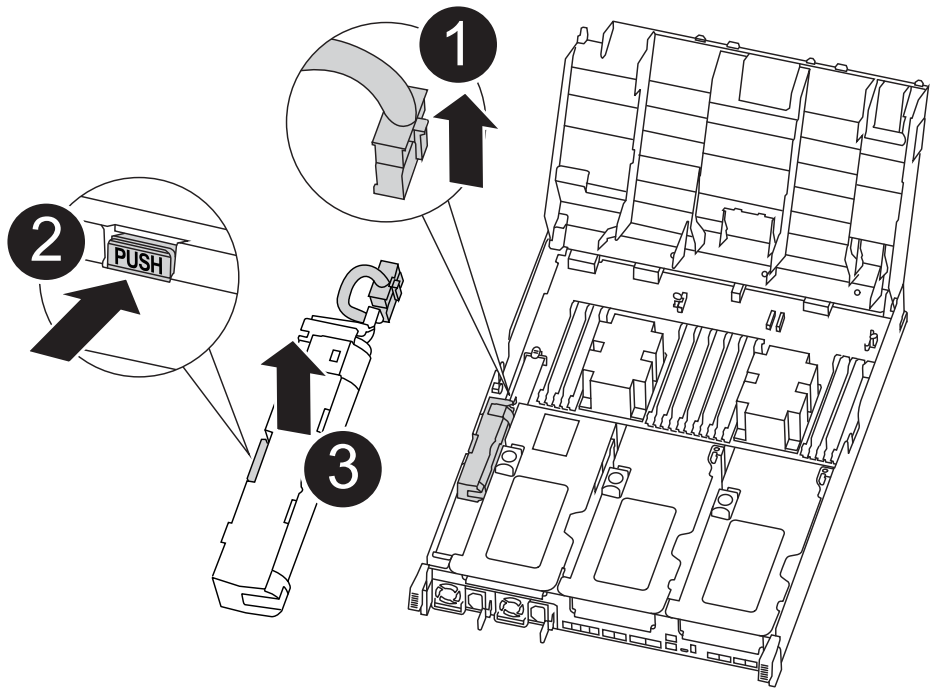

1	Battery plug
2	Locking tab
3	NVDIMM battery

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Remove the replacement battery from its package.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.
7. Plug the battery plug back into the controller module, and then close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module

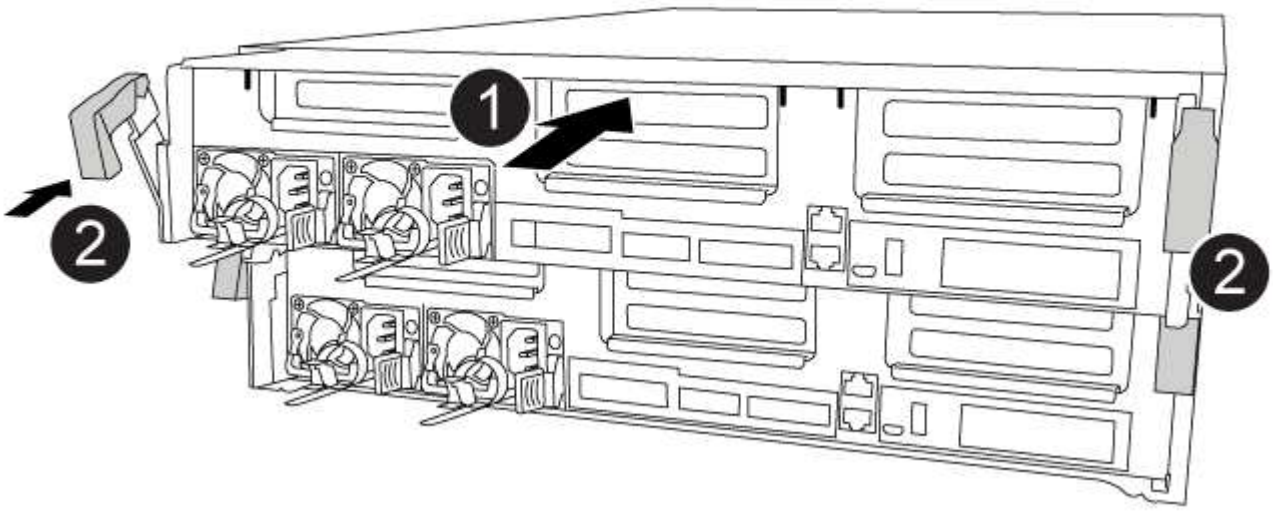

1	Controller module
2	Controller locking latches

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

> Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an NVDIMM - AFF A400

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animations, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

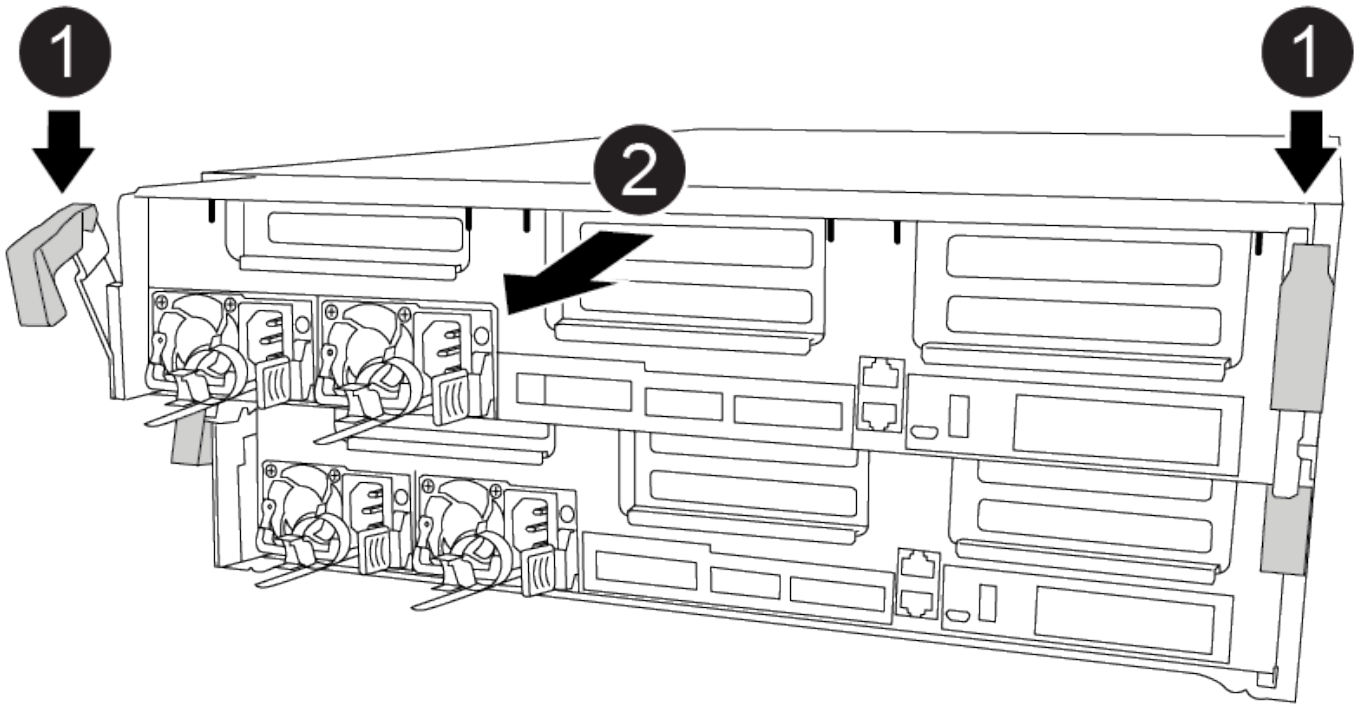

1	Locking latches
2	Controller moves slightly out of chassis

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the FRU map on top of the air duct or the FRU Map on the top of the slot 1 riser.

- The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.
- Although the contents of the NVDIMM is encrypted, it is a best practice to erase the contents of the NVDIMM before replacing it. For more information, see the Statement of Volatility on the NetApp Support Site.

> You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the NVDIMM.

> The animation shows empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace the NVDIMM

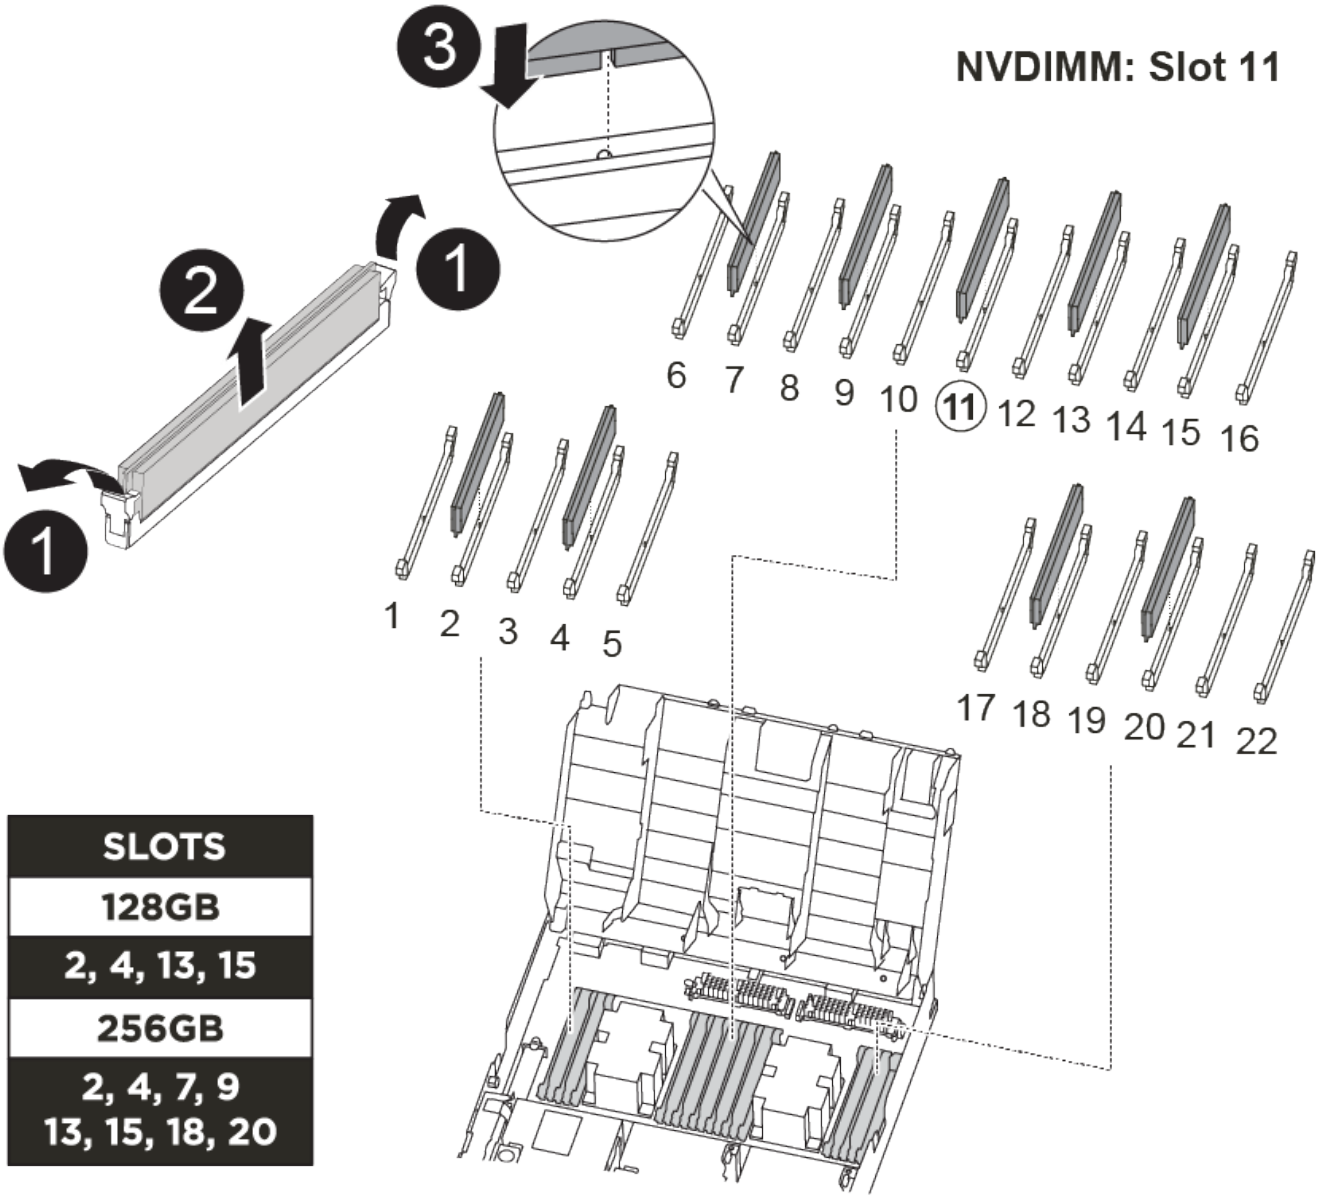

1	DIMM locking tabs
2	DIMM
3	DIMM socket

1. Open the air duct and then locate the NVDIMM in slot 11 on your controller module.

   The NVDIMM looks significantly different than system DIMMs.

2. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

3. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

4. Locate the slot where you are installing the NVDIMM.
5. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
7. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot

process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`

4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    -------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    -------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe or mezzanine card - AFF A400

To replace a PCIe or mezzanine card, you must disconnect the cables and any SFP and QSFP modules from the cards, replace the failed PCIe or mezzanine card, and then recable the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

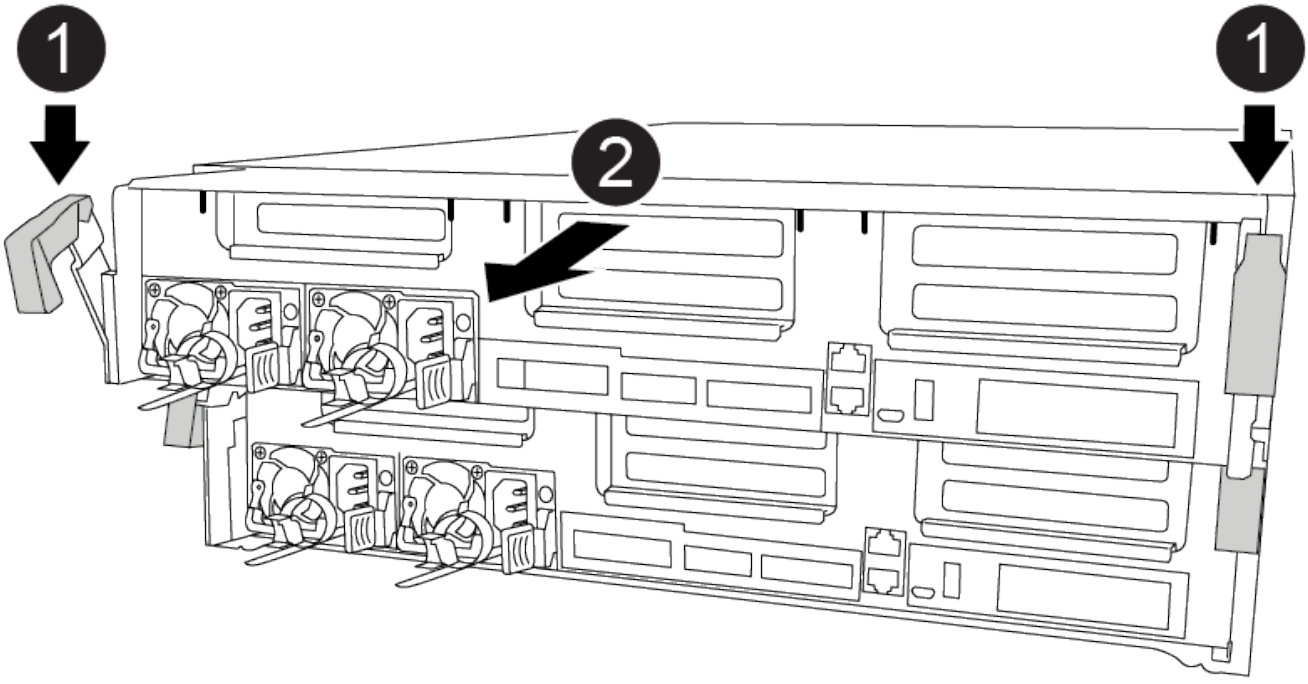

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace a PCIe card

To replace a PCIe card, you must locate the failed PCIe card, remove the riser that contains the card from the controller module, replace the card, and then reinstall the PCIe riser in the controller module.

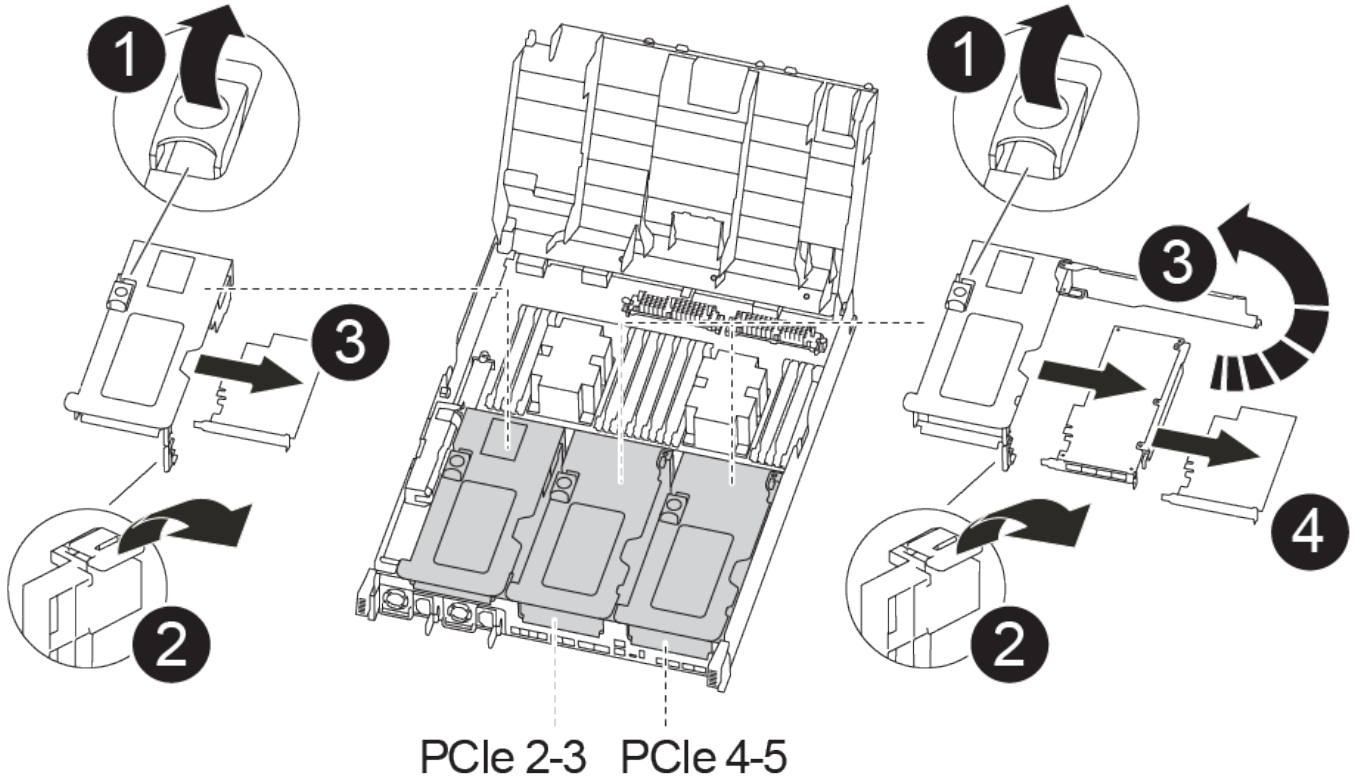

1	Riser locking latch
2	PCI card locking latch
3	PCI locking plate
4	PCI card

1. Remove the riser containing the card to be replaced:
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up straight up and set it aside on a stable flat surface,
2. Remove the PCIe card from the riser:

a. Turn the riser so that you can access the PCIe card.
b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
c. For risers 2 and 3 only, swing the side panel up.
d. Remove the PCIe card from the riser by gently pushing up on the bracket and lift the card straight out of the socket.

3. Install the replacement PCIe card in the riser by aligning the card with the socket, press the card into the socket and then close the side panel on the riser, if present.

   Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The PCIe card must be fully and evenly seated in the slot.

   > If you are installing a card in the bottom slot and cannot see the card socket well, remove the top card so that you can see the card socket, install the card, and then reinstall the card you removed from the top slot.

4. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 4: Replace the mezzanine card

The mezzanine card is located under riser number 3 (slots 4 and 5). You must remove that riser to access the mezzanine card, replace the mezzanine card, and then reinstall riser number 3. See the FRU map on the controller module for more information.

You can use the following animation, illustration, or the written steps to replace the mezzanine card.

Animation - Replace the mezzanine card

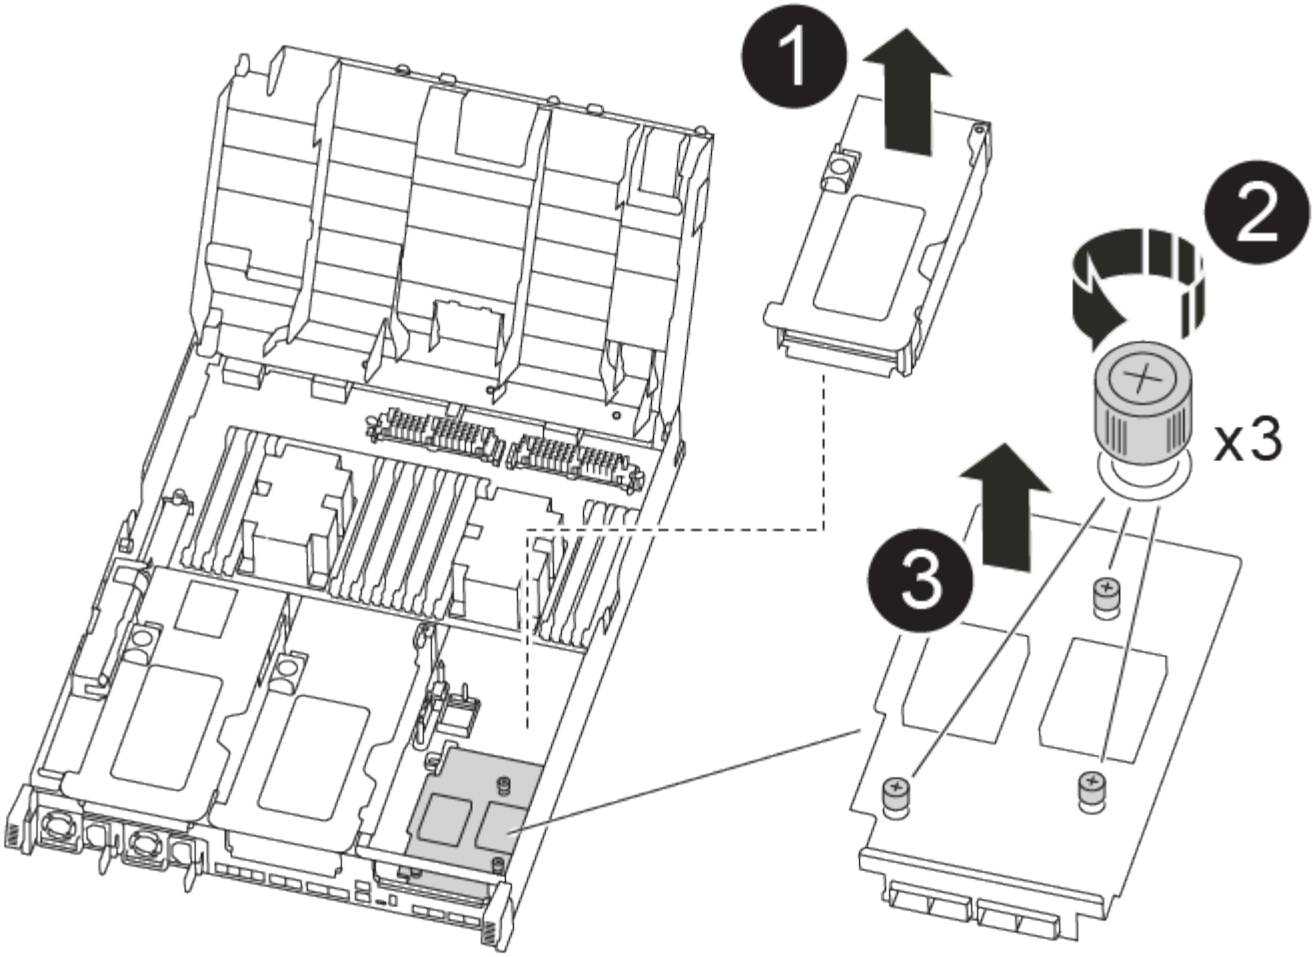

1	PCI riser
2	Riser thumbscrew
3	Riser card

1. Remove riser number 3 (slots 4 and 5):
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up, and then set it aside on a stable, flat surface.
2. Replace the mezzanine card:
   a. Remove any QSFP or SFP modules from the card.
   b. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket and

set it aside.

c. Align the replacement mezzanine card over the socket and the guide pins and gently push the card into the socket.

d. Tighten the thumbscrews on the mezzanine card.

3. Reinstall the riser:

a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.

b. Push the riser squarely into the socket on the motherboard.

c. Rotate the latch down flush with the sheet metal on the riser.

Step 5: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the installation of the controller module:

a. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.

Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

5. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
6. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Restore the controller module to operation

To restore the controller, you must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 7: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the

surviving cluster.

5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 8: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replacing a power supply - AFF A400

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting the replacement PSU to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

You can use the following illustration with the written steps to replace the power supply.

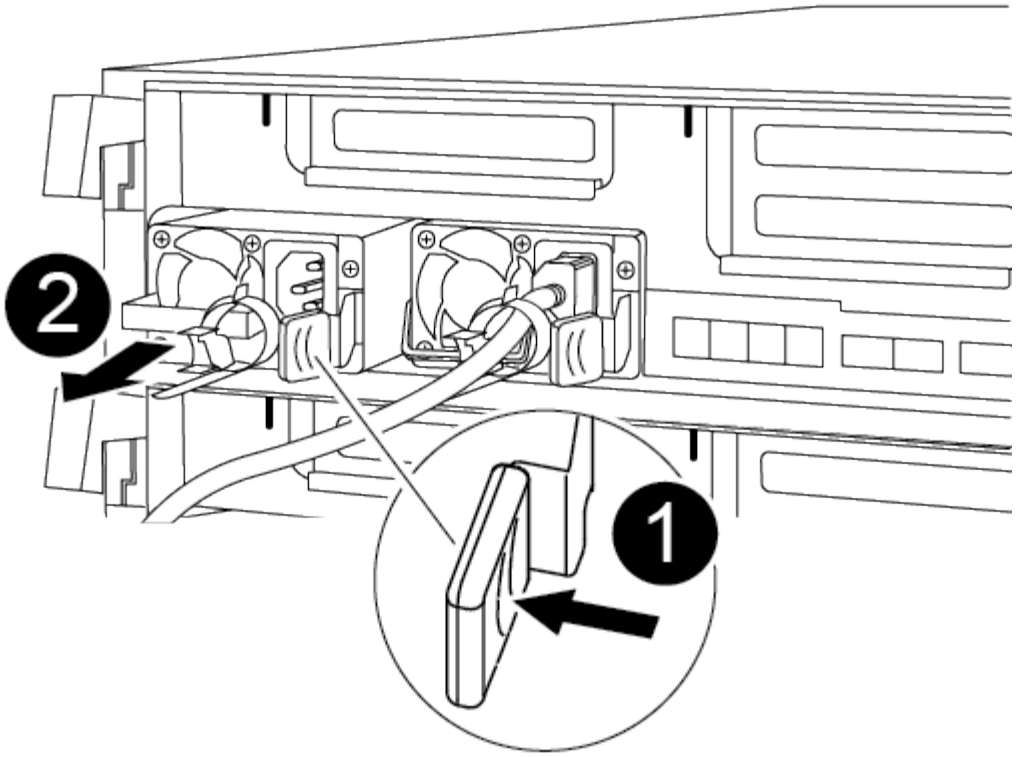

1	PSU locking tab
2	Power cable retainer

1. If you are not already grounded, properly ground yourself.
2. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
3. Disconnect the power supply:
   a. Open the power cable retainer, and then unplug the power cable from the power supply.
   b. Unplug the power cable from the power source.
4. Remove the power supply:
   a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
   b. Press the blue locking tab to release the power supply from the chassis.
   c. Using both hands, pull the power supply out of the chassis, and then set it aside.
5. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

6. Rotate the cam handle so that it is flush against the power supply.
7. Reconnect the power supply cabling:
    a. Reconnect the power cable to the power supply and the power source.
    b. Secure the power cable to the power supply using the power cable retainer.

    Once power is restored to the power supply, the status LED should be green.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A400

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

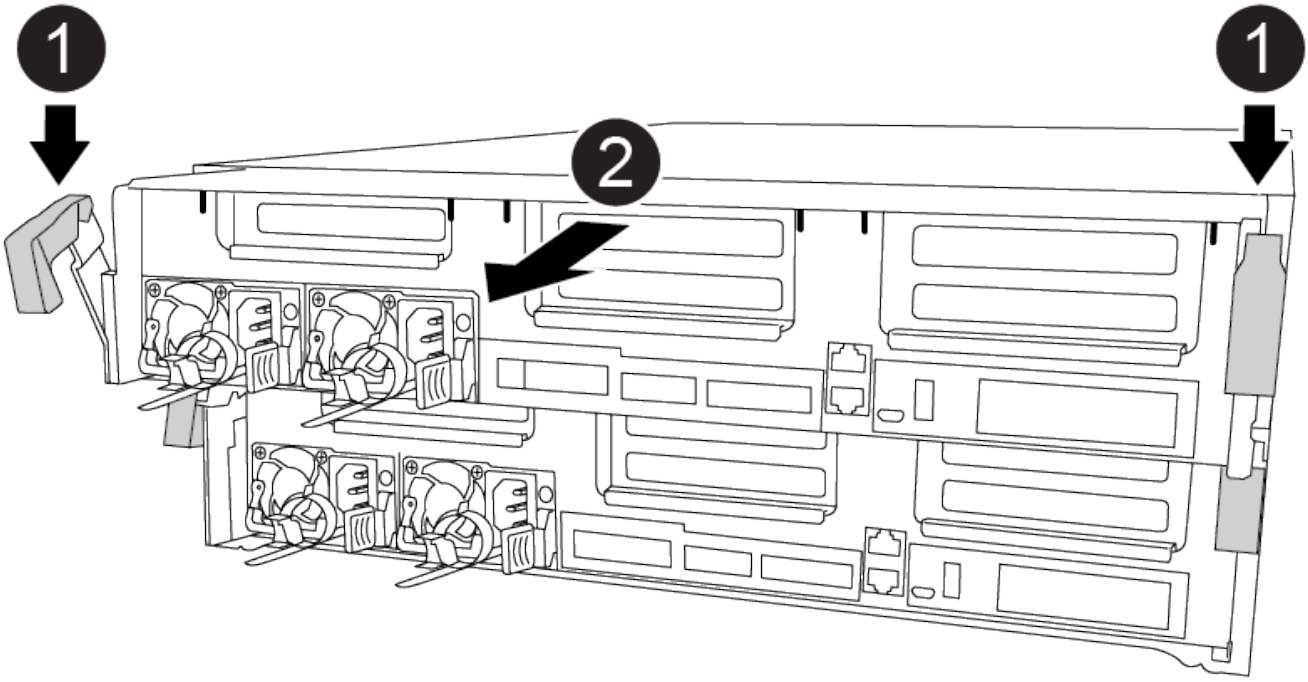

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the RTC battery

You need to locate the RTC battery inside the controller module, and then follow the specific sequence of steps. See the FRU map inside the controller module for the location of the RTC battery.

You can use the following animation, illustration, or the written steps to replace the RTC battery.

Animation - Replace the RTC battery

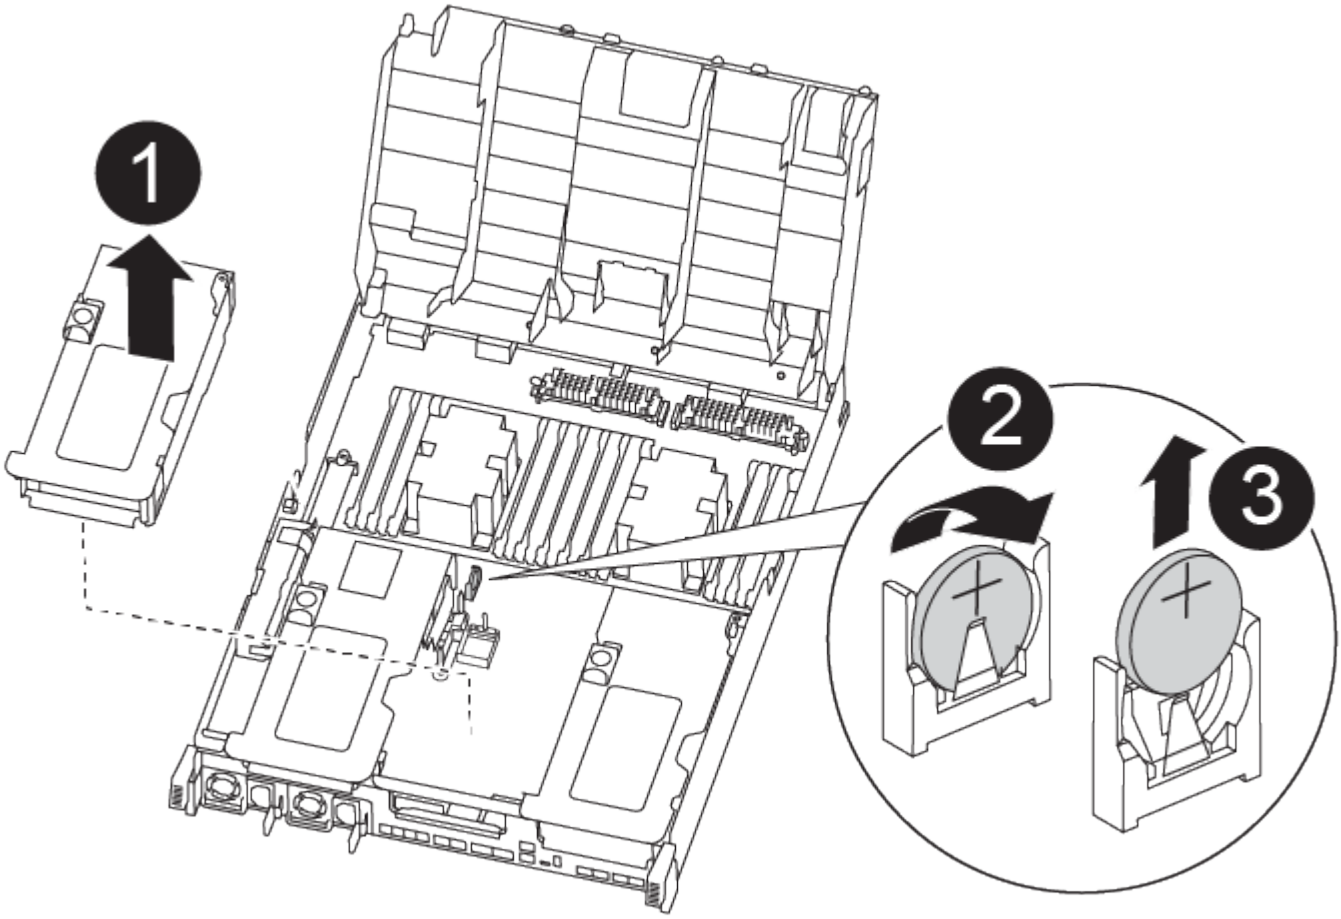

1	Middle riser
2	Remove RTC battery
3	Seat RTC battery

1. If you are not already grounded, properly ground yourself.
2. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

3. Locate, remove, and then replace the RTC battery:
   a. Using the FRU map, locate the RTC battery on the controller module.
   b. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

   c. Remove the replacement battery from the antistatic shipping bag.
   d. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
5. Close the air duct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module

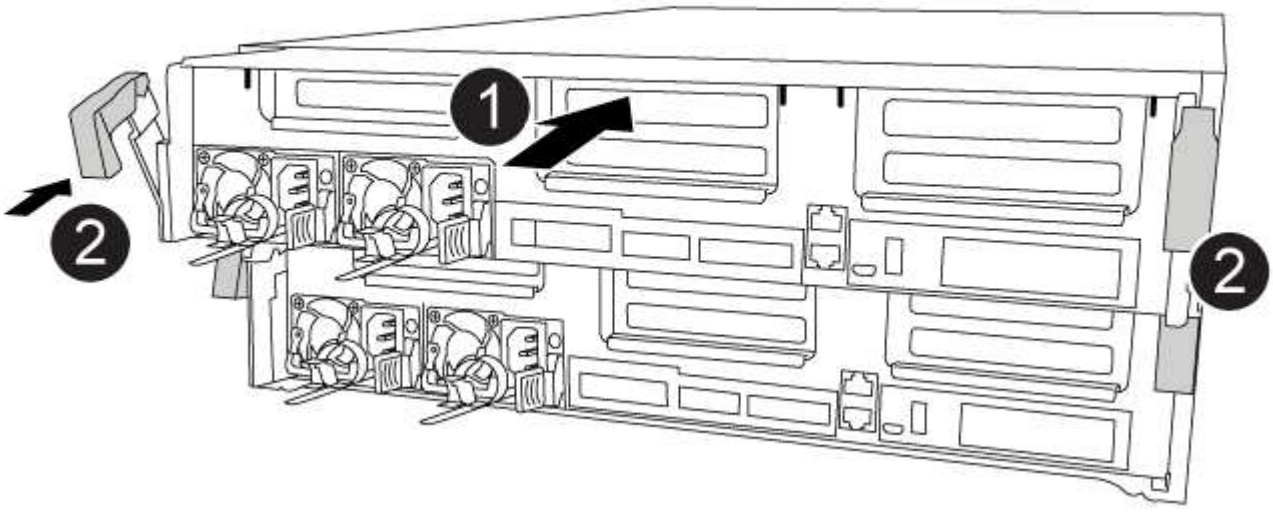

1	Controller module
2	Controller locking latches

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      > If your system stops at the boot menu, select the option to boot to LOADER.

5. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF A800 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations (including ASA configurations), you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick steps - AFF A800

This guide gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use the AFF A800 Installation and Setup Instructions if you are familiar with installing NetApp systems.

Video steps - AFF A800

The following video shows how to install and cable your new system.

"Animation - Installation and Setup of an AFF A800

Detailed steps - AFF A800

This section gives detailed step-by-step instructions for installing an AFF A800 system.

Step 1: Prepare for installation

To install your AFF A800 system, you need to create an account and register the system. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the NetApp Hardware Universe (HWU) for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

What you need

You need to provide the following at your site:

- Rack space for the storage system
  - 4U in an HA configuration for the platform
  - 2U for each NS224 storage shelf
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

Steps

1. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register (NetApp Product Registration) your system.
2. Download and install NetApp Downloads: Config Advisor on your laptop.
3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Connector type	Part number and length	Type of cable…	For…
100 GbE cable	X66211A-05 (112-00595), 0.5m X66211-1 (112-00573), 1m X66211-2 (112-00574), 2m X66211-5 (112-00576), 5m		• HA interconnect • Cluster interconnect network • Storage, Data
10 GbE cable	X6566B-3-R6 (112-00300), 3m; X6566B-5-R6 (112-00301), 5m		• Data

Connector type	Part number and length	Type of cable…	For…
25 GbE cable	X66240A-2 (112-00598), 2m; X66240A-5 (112-00600), 5m		• Data
RJ-45 (order dependent)	Not applicable		• Management
Fibre Channel	X66250-2 (112-00342) 2m; X66250-5 (112-00344) 5m; X66250-15 (112-00346) 15m; X66250-30 (112-00347) 30m		• Network
Micro-USB console cable	Not applicable		• Console connection during software setup
Power cables	Not applicable		Connecting the PSUs to power source

4. Download and complete the Cluster Configuration Worksheet.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.

   Installing SuperRail into a four-post rack

2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

3. Attach cable management devices (as shown).

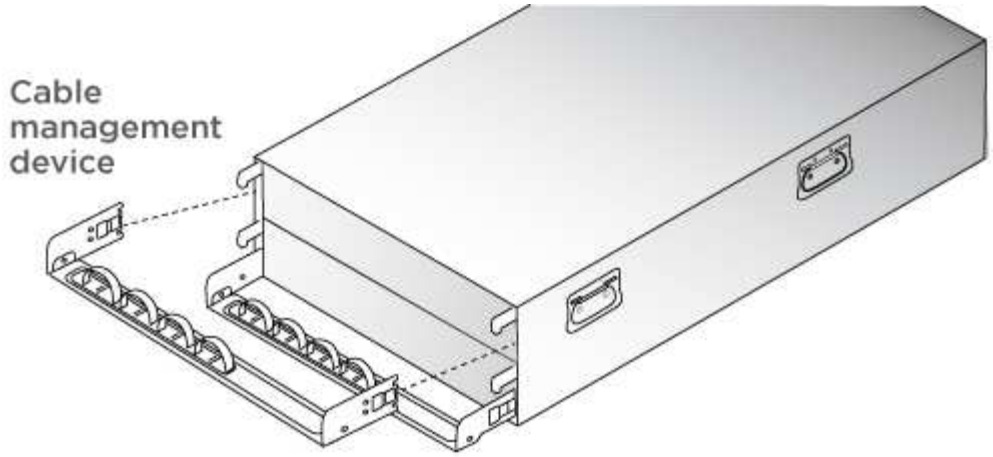

4. Place the bezel on the front of the system.

Step 3: Cable controllers

There is required cabling for your platform's cluster using the two-node switchless cluster method or the cluster interconnect network method. There is optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cable to a host network and storage.

Required cabling: Cable controllers to a cluster

Cable the controllers to a cluster by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Cable a two-node switchless cluster

Management network ports on the controllers are connected to switches. The HA interconnect and cluster interconnect ports are cabled on both controllers.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

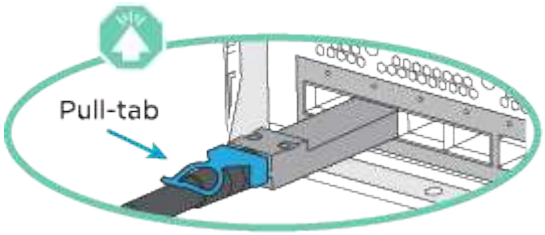

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or the tabulated steps to complete the cabling between the controllers and the switches:

   Animation - Cable a two-node switchless cluster

Step	Perform on each controller module
1	Cable the HA interconnect ports: • e0b to e0b • e1b to e1b 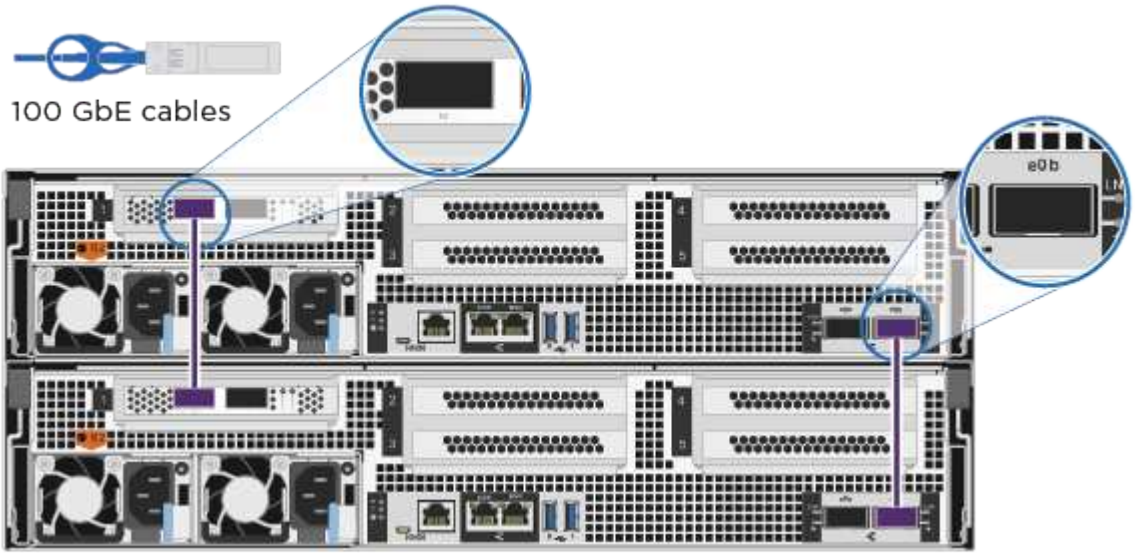
2	Cable the cluster interconnect ports: • e0a to e0a • e1a to e1a 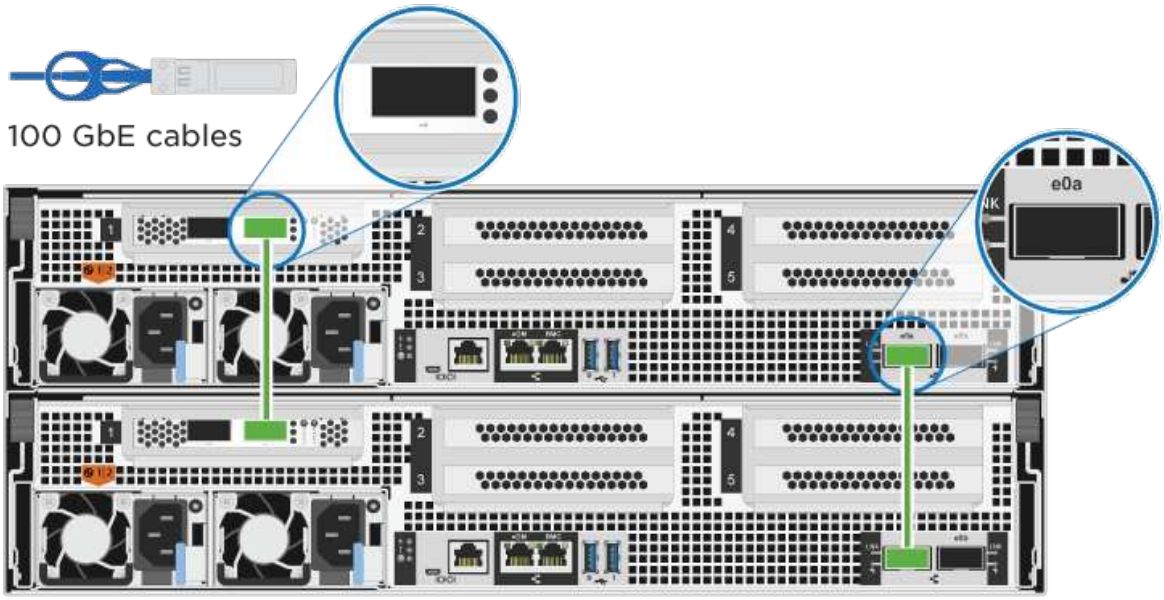

Step	Perform on each controller module
3	Cable the management ports to the management network switches 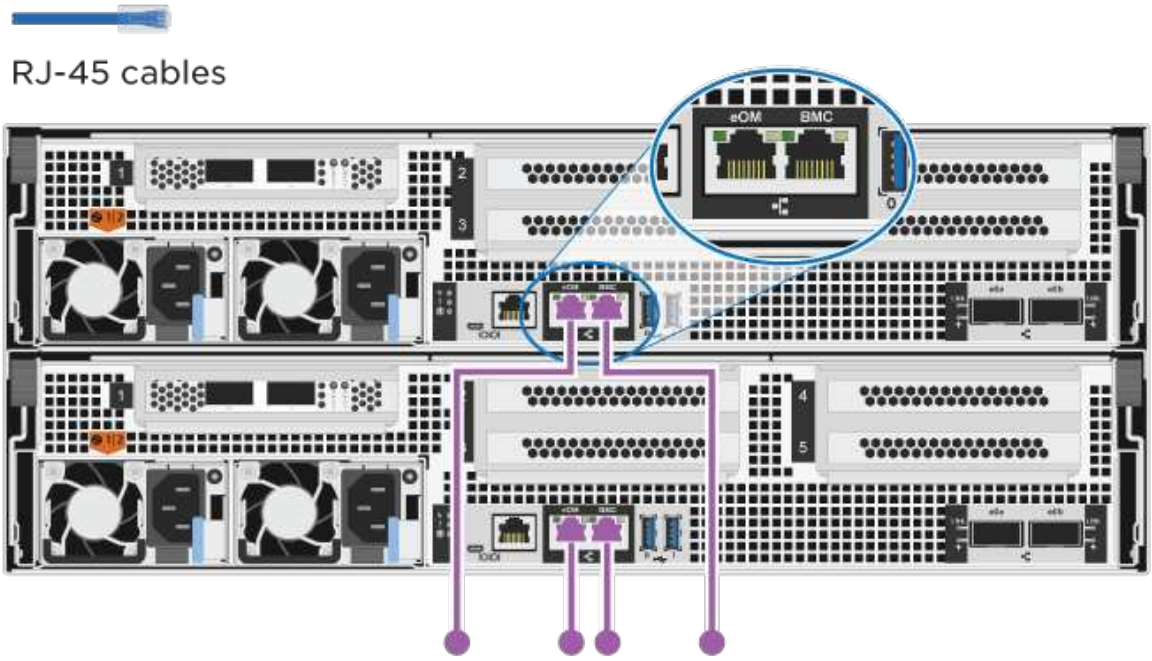
	DO NOT plug in the power cords at this point.

2. To perform optional cabling, see:
   - Option 1: Cable to a Fibre Channel host network
   - Option 2: Cable to a 10GbE host network
   - Option 3: Cable the controllers to a single drive shelf
   - Option 4: Cable the controllers to two drive shelves
3. To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 2: Cable a switched cluster

Cluster interconnect and management network ports on the controllers are connected to switches while the HA interconnect ports are cabled on both controllers.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

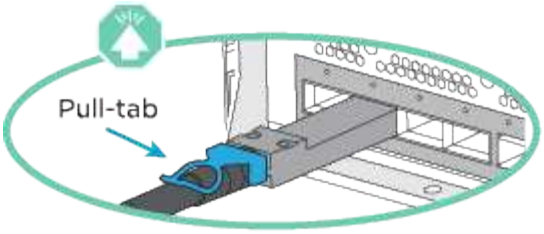

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or the tabulated steps to complete the cabling between the controllers and the switches:

Animation - Cable a switched cluster

Step	Perform on each controller module
1	Cable the HA interconnect ports: • e0b to e0b • e1b to e1b 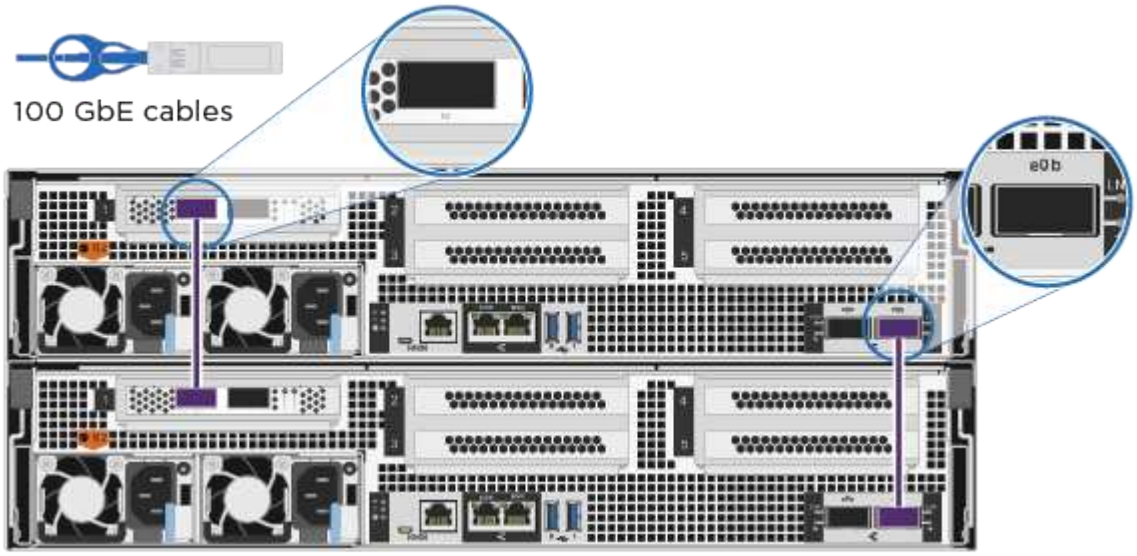

Step	Perform on each controller module
2	Cable the cluster interconnect ports to the 100 GbE cluster interconnect switches. • e0a • e1a 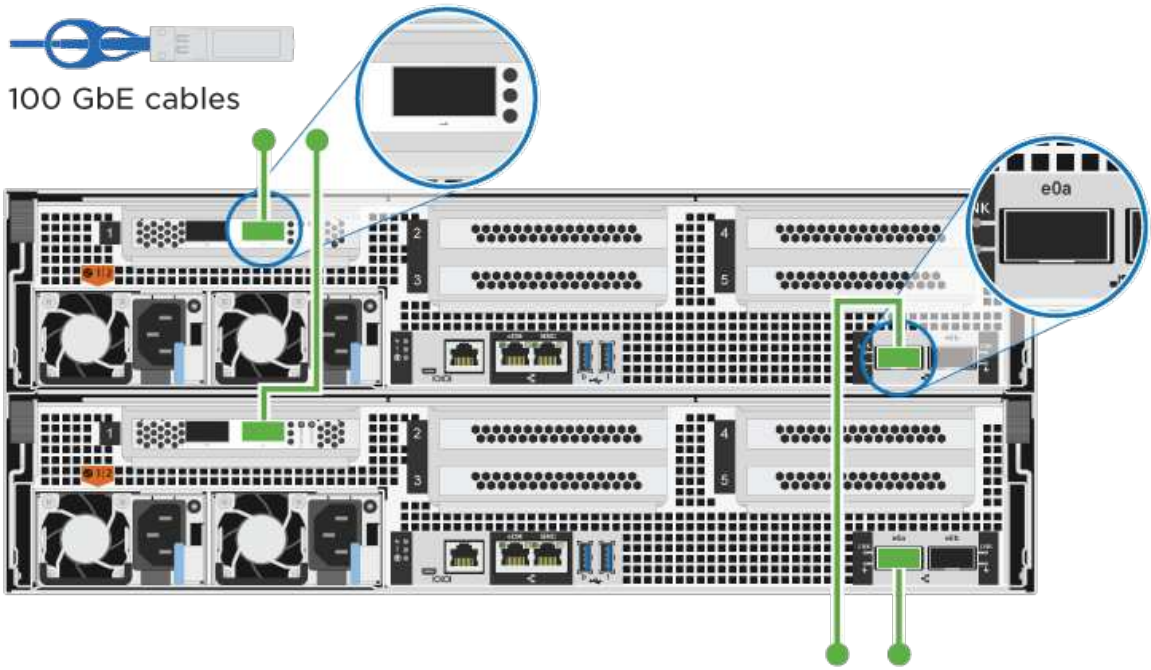
3	Cable the management ports to the management network switches 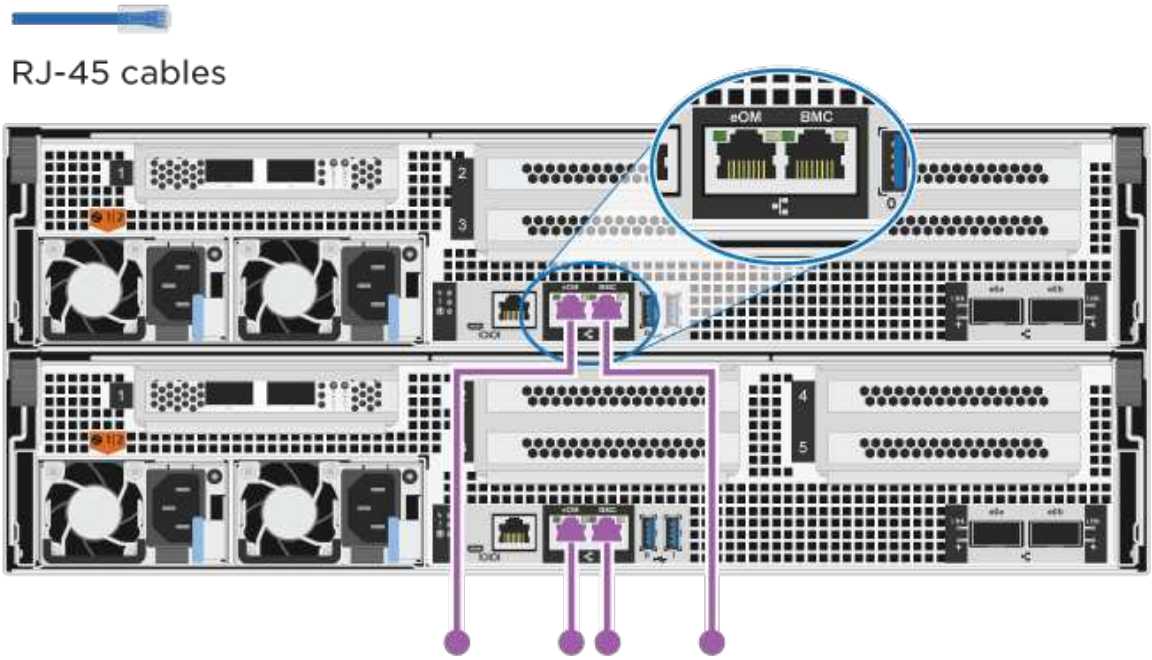

Step	Perform on each controller module
	DO NOT plug in the power cords at this point.

2. To perform optional cabling, see:
   - Option 1: Cable to a Fibre Channel host network
   - Option 2: Cable to a 10GbE host network
   - Option 3: Cable the controllers to a single drive shelf
   - Option 4: Cable the controllers to two drive shelves
3. To complete setting up your system, see Step 4: Complete system setup and configuration.

Optional cabling: Cable configuration-dependent options

You have configuration-dependent optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cabling to a host network and storage.

Option 1: Cable to a Fibre Channel host network

Fibre Channel ports on the controllers are connected to Fibre Channel host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

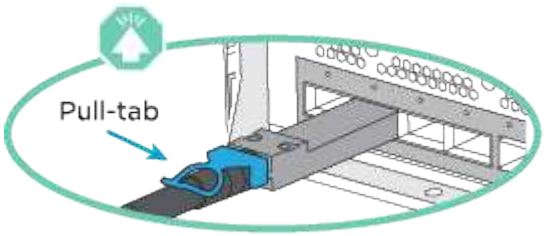

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Step	Perform on each controller module
1	Cable ports 2a through 2d to the FC host switches. 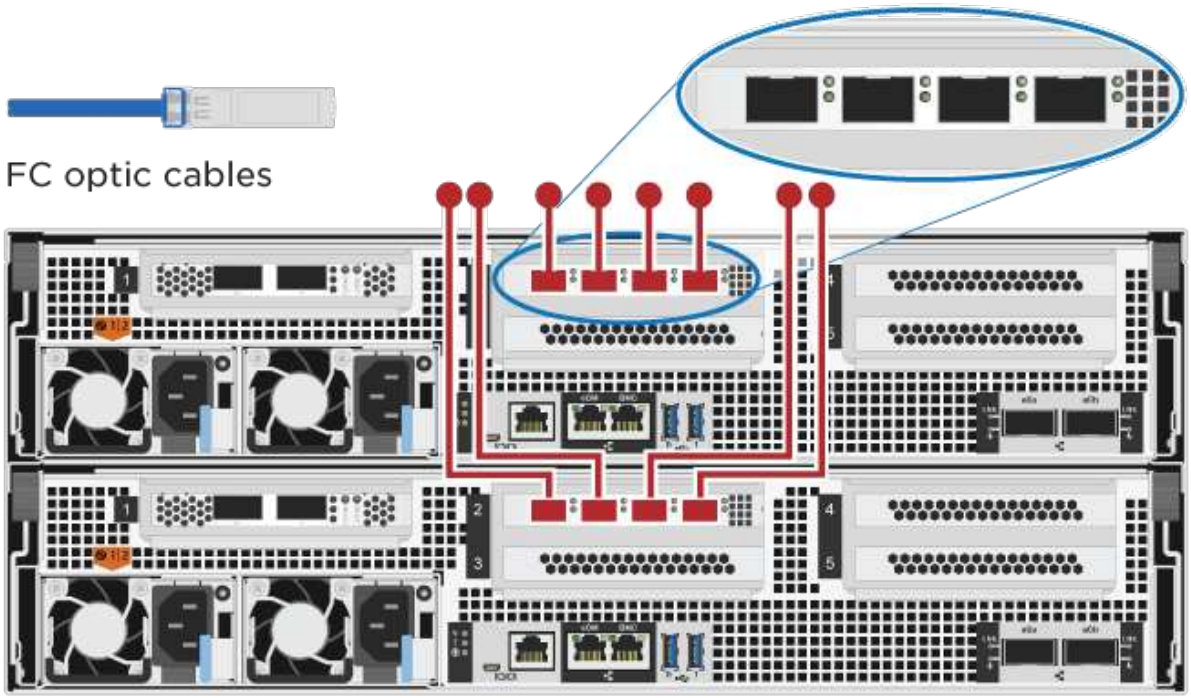
2	To perform other optional cabling, choose from: • Option 3: Cable the controllers to a single drive shelf • Option 4: Cable the controllers to two drive shelves
3	To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 2: Cable to a 10GbE host network

10GbE ports on the controllers are connected to 10GbE host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

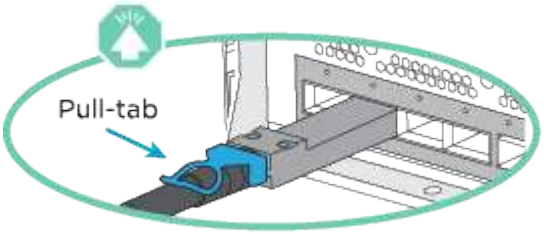

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Step	Perform on each controller module
1	Cable ports e4a through e4d to the 10GbE host network switches. 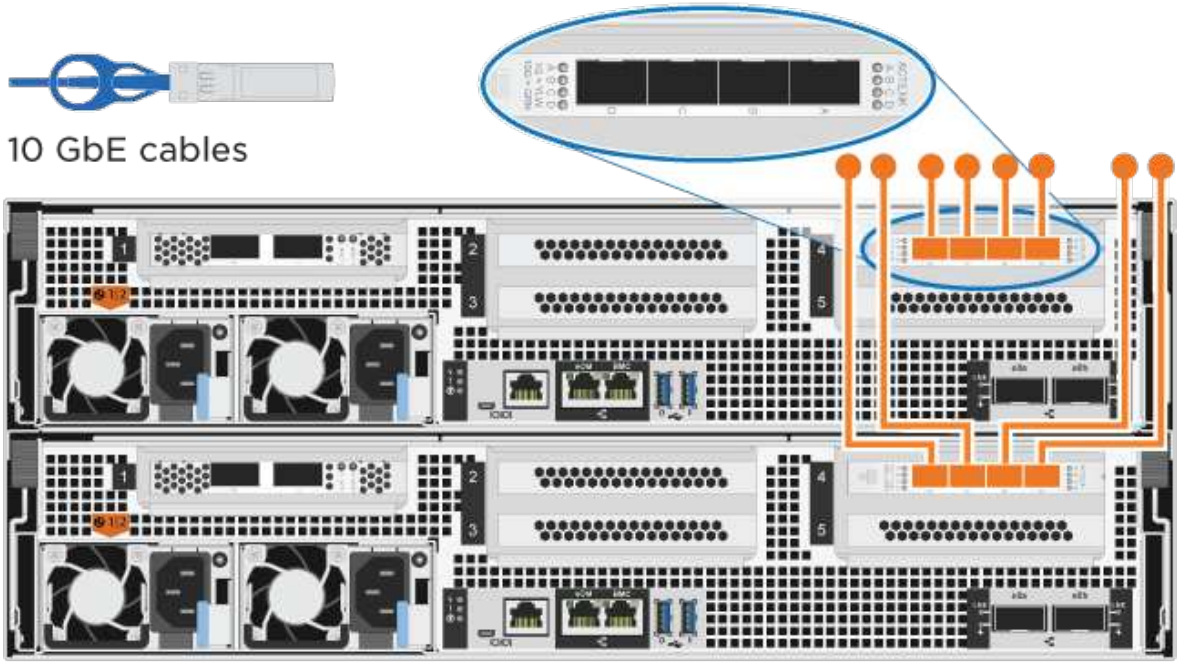
2	To perform other optional cabling, choose from: • Option 3: Cable the controllers to a single drive shelf • Option 4: Cable the controllers to two drive shelves
3	To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 3: Cable the controllers to a single drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

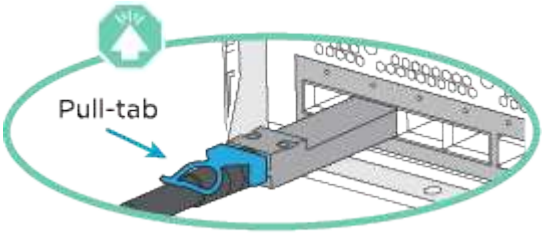

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to cable your controllers to a single shelf:

Animation - Cable the controllers to a single drive shelf

Step	Perform on each controller module
1	Cable controller A to the shelf: 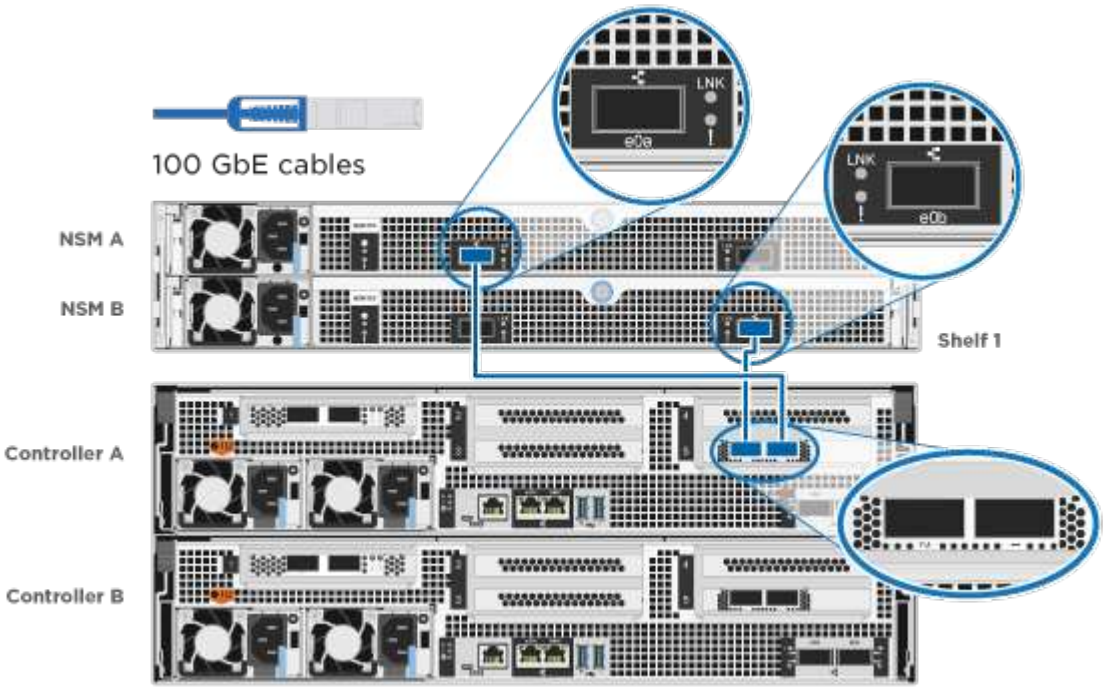
2	Cable controller B to the shelf: 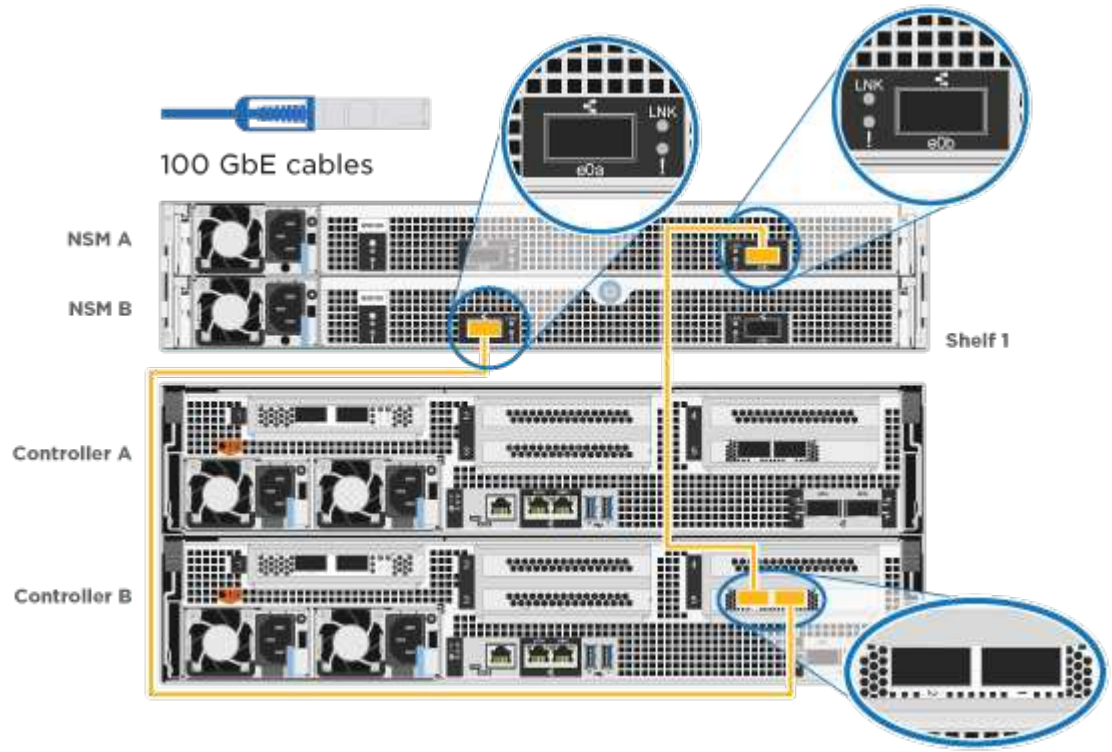

To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 4: Cable the controllers to two drive shelves

You must cable each controller to the NSM modules on both NS224 drive shelves.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

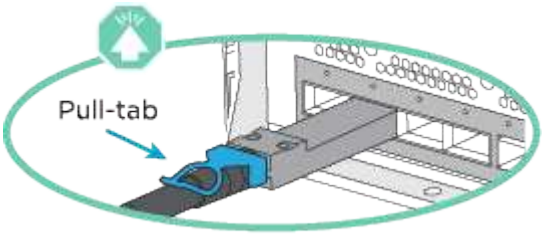

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to cable your controllers to two drive shelves:

Animation - Cable the controllers to two drive shelves

Step	Perform on each controller module
1	Cable controller A to the shelves: 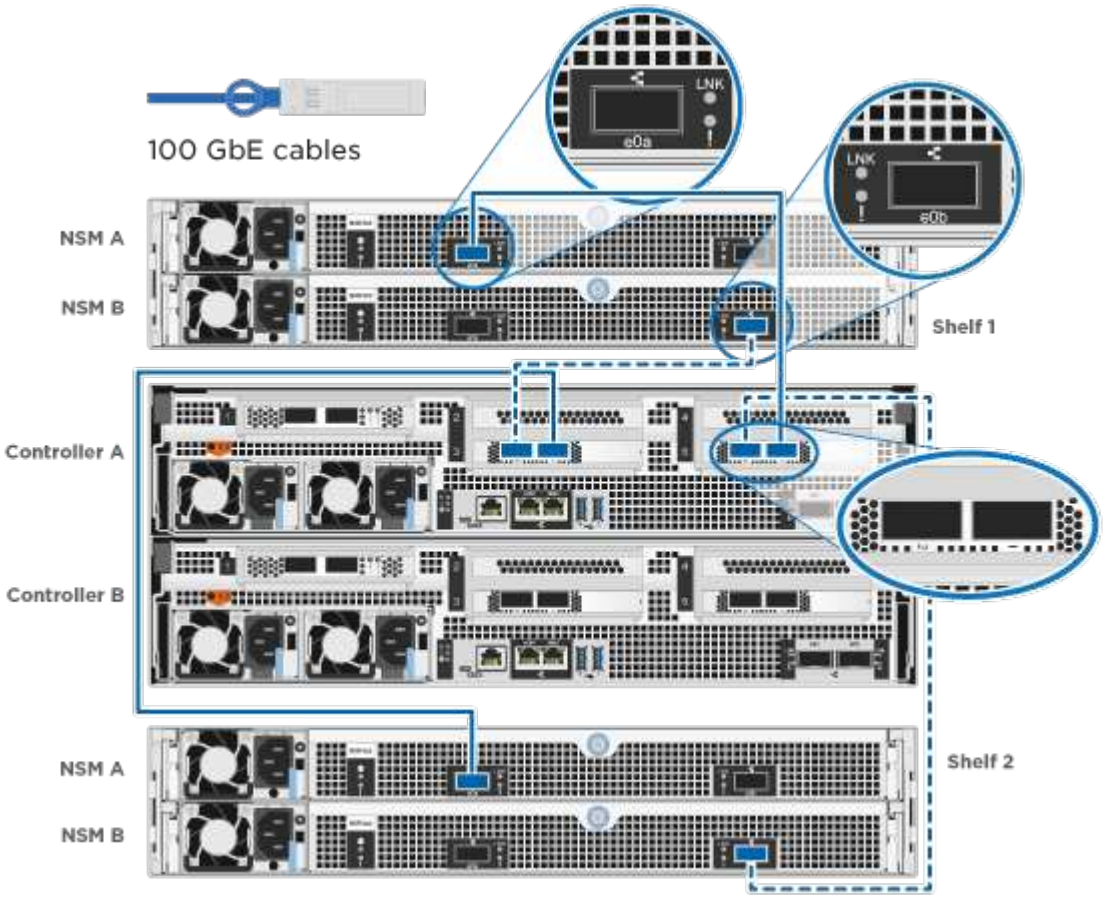

Step	Perform on each controller module
2	Cable controller B to the shelves: 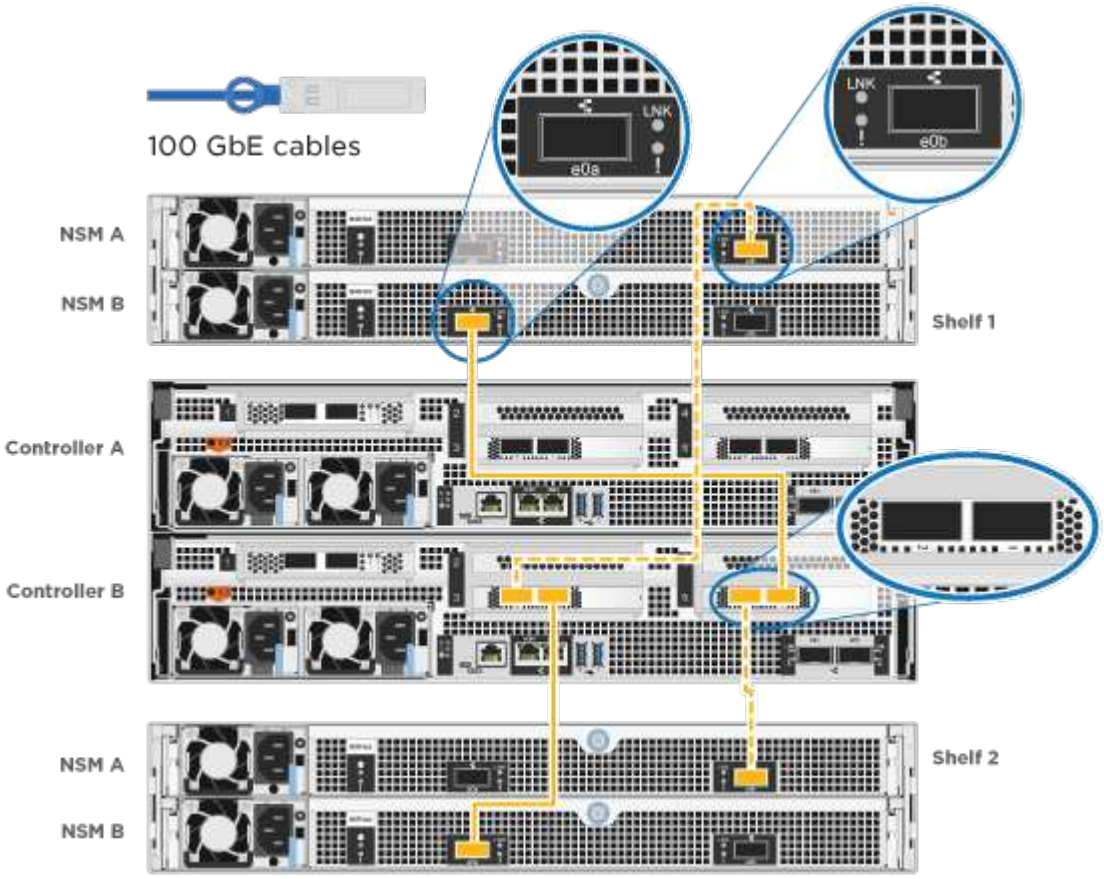

To complete setting up your system, see Step 4: Complete system setup and configuration.

Step 4: Complete system setup and configuration

Complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Complete system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

2. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

3. Use the animation to connect your laptop to the Management switch:

Animation - Connect your laptop to the Management switch

4. Select an ONTAP icon listed to discover:

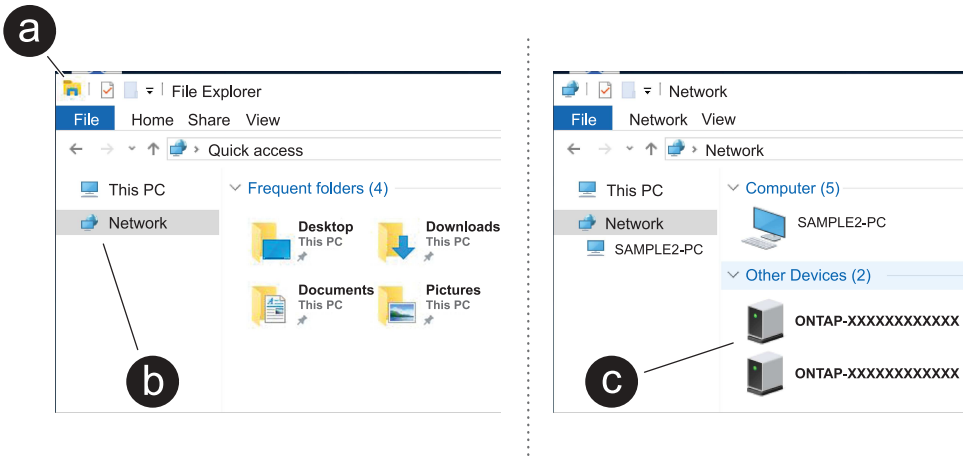

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

      System Manager opens.

5. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
6. Verify the health of your system by running Config Advisor.
7. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Complete system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your system.

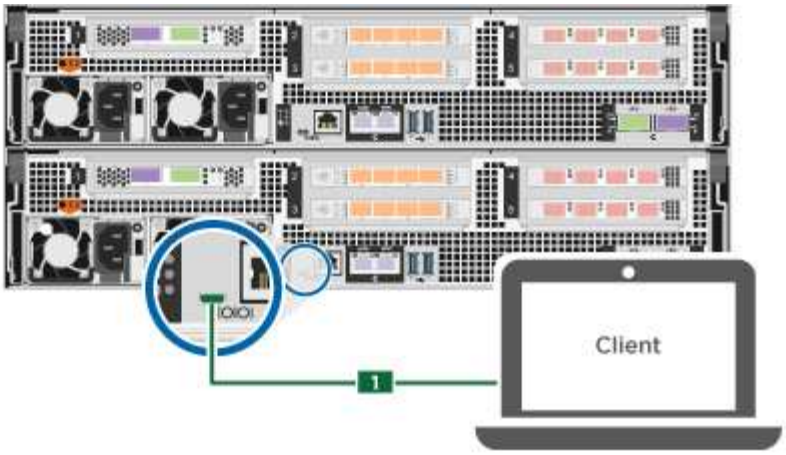

c. Connect the laptop or console to the switch on the management subnet.

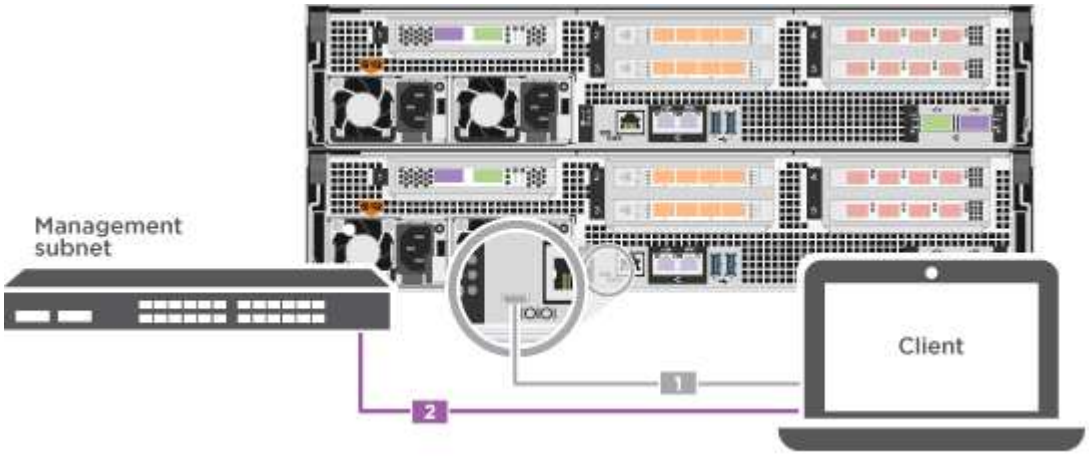

d. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

3. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

4. Using System Manager on your laptop or console, configure your cluster:

a. Point your browser to the node management IP address.

> The format for the address is https://x.x.x.x.

b. Configure the system using the data you collected in the ONTAP Configuration Guide.

5. Verify the health of your system by running Config Advisor.
6. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain AFF A800 hardware

Maintain the hardware of your AFF A800 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the AFF A800 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF A800 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Drive	A drive is a device that provides the physical storage media for data.

Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.
NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.
PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - AFF A800

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your AFF A800 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5 **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - AFF A800

Before replacing the boot media in your AFF A800, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - AFF A800

Shut down the impaired controller in your AFF A800 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - AFF A800

The boot media in your AFF A800 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module,

removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

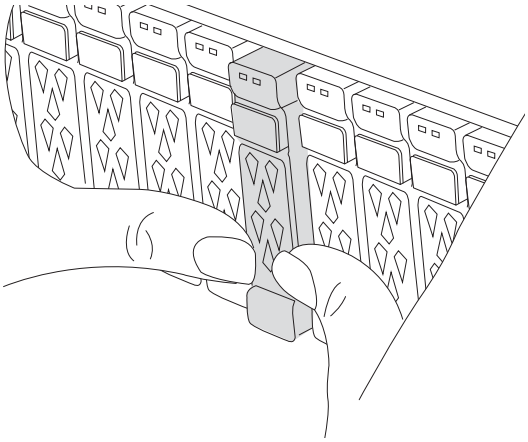

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

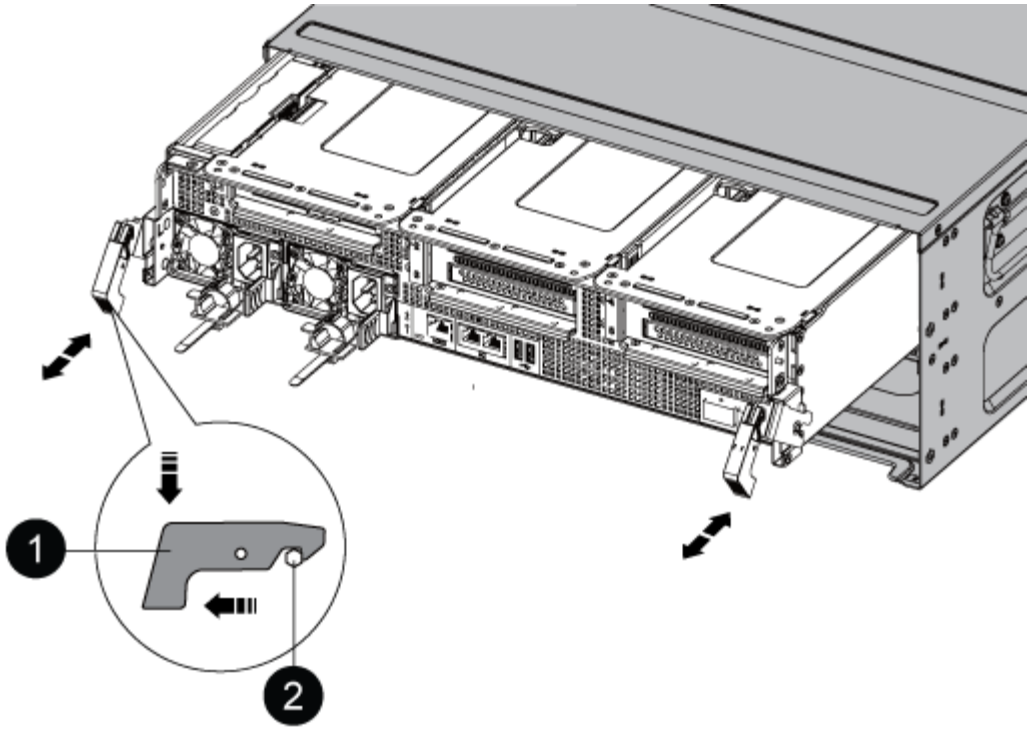

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

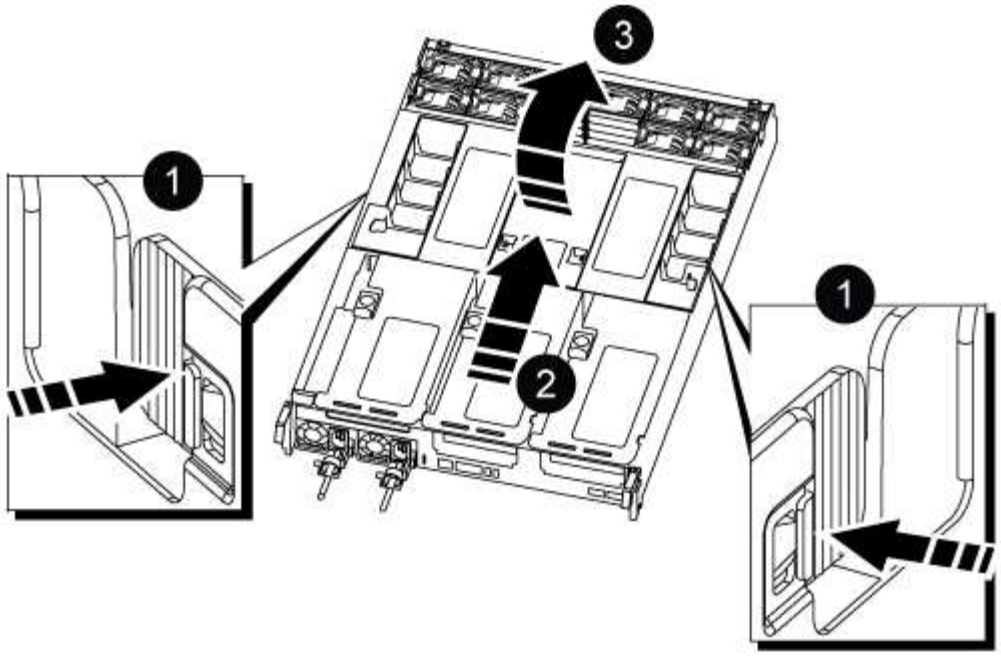

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

10. Locate the boot media in the controller module and replace it:

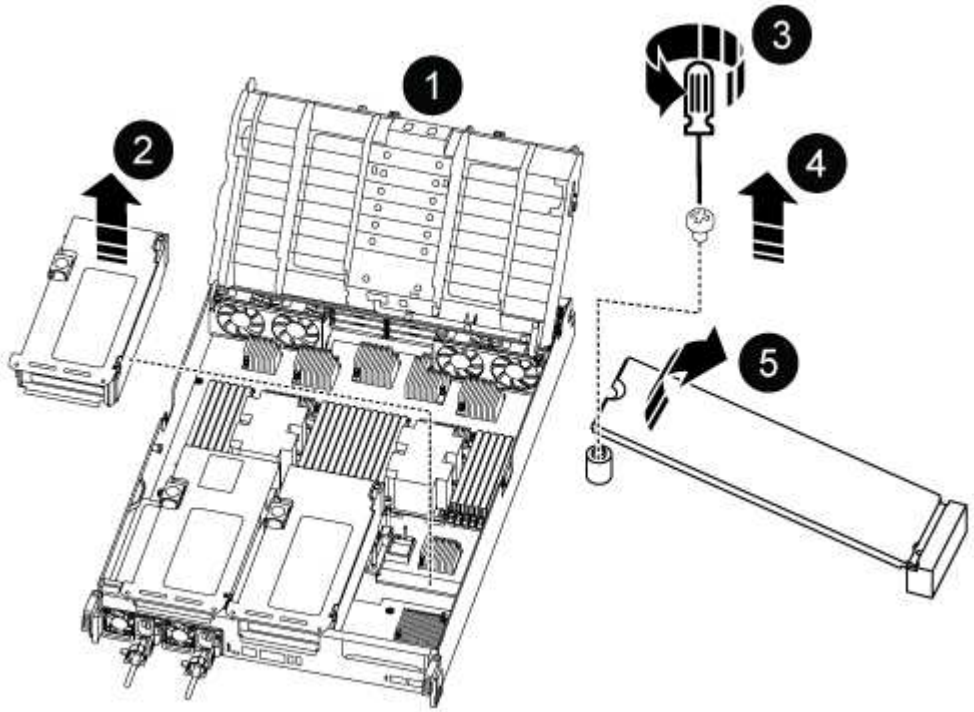

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.

b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.

11. Install the replacement boot media into the controller module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down toward the motherboard.

c. Secure the boot media to the motherboard using the boot media screw.

Do not over-tighten the screw or you might damage the boot media.

12. Reinstall the riser into the controller module.

13. Close the air duct:

a. Rotate the air duct downward.

b. Slide the air duct toward the risers until it clicks into place.

14. Install the controller module:

a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module half-way into the way into the system.

b. Recable the controller module, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

The controller module begins to boot and stops at the LOADER prompt.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - AFF A800

After installing the new boot media device in your AFF A800 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and

determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

   Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

   Show example of configuration error finding prompts

   ```
   Error when fetching key manager config from partner ${partner_ip}:
   ${status}

   Has key manager been configured on this system

   Is the key manager onboard
   ```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - AFF A800

If a component in your AFF A800 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF A800

Get started with replacing the boot media in your AFF A800 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - AFF A800

Before replacing the boot media in your AFF A800 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - AFF A800

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - AFF A800

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: System is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media and prepare for manual boot recovery - AFF A800

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

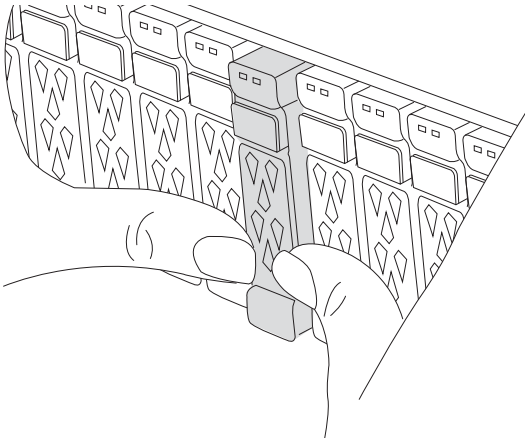

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management

device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

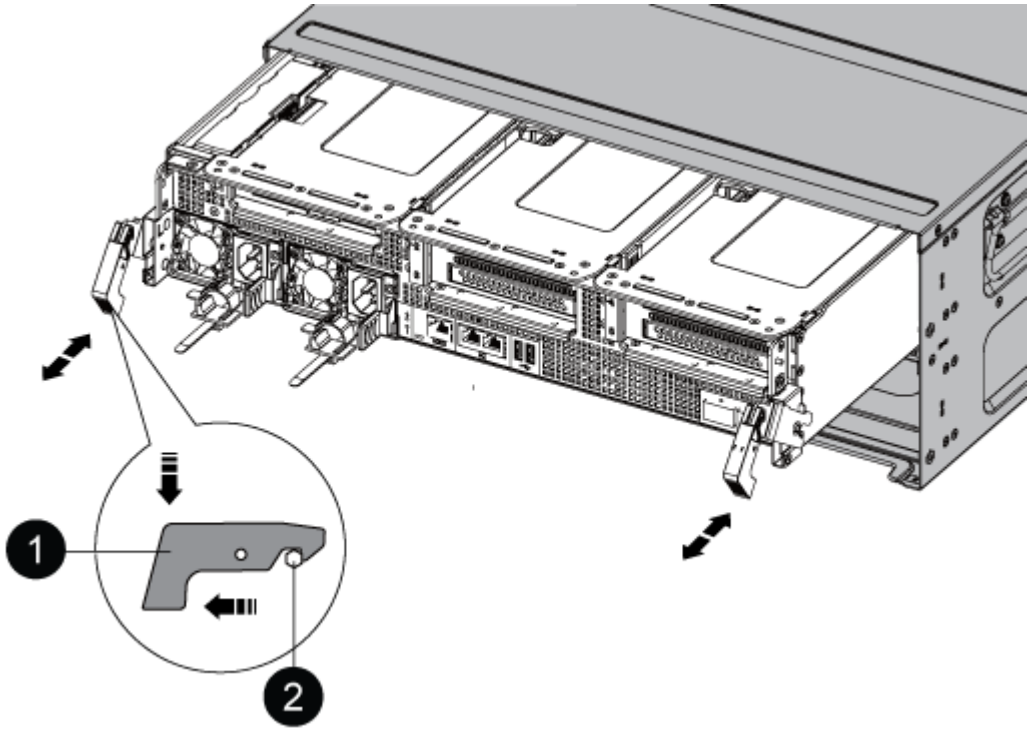

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

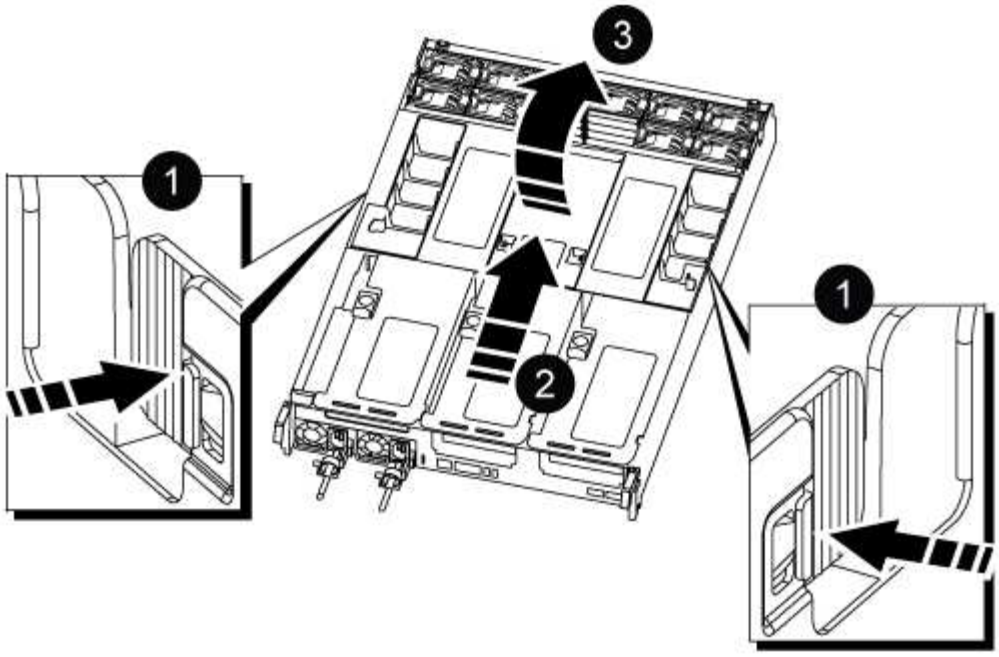

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 2: Replace the boot media

You locate the failed boot media in the controller module by removing Riser 3 on the controller module before you can replace the boot media.

You need a Phillips head screwdriver to remove the screw that holds the boot media in place.

1. Locate the boot media:

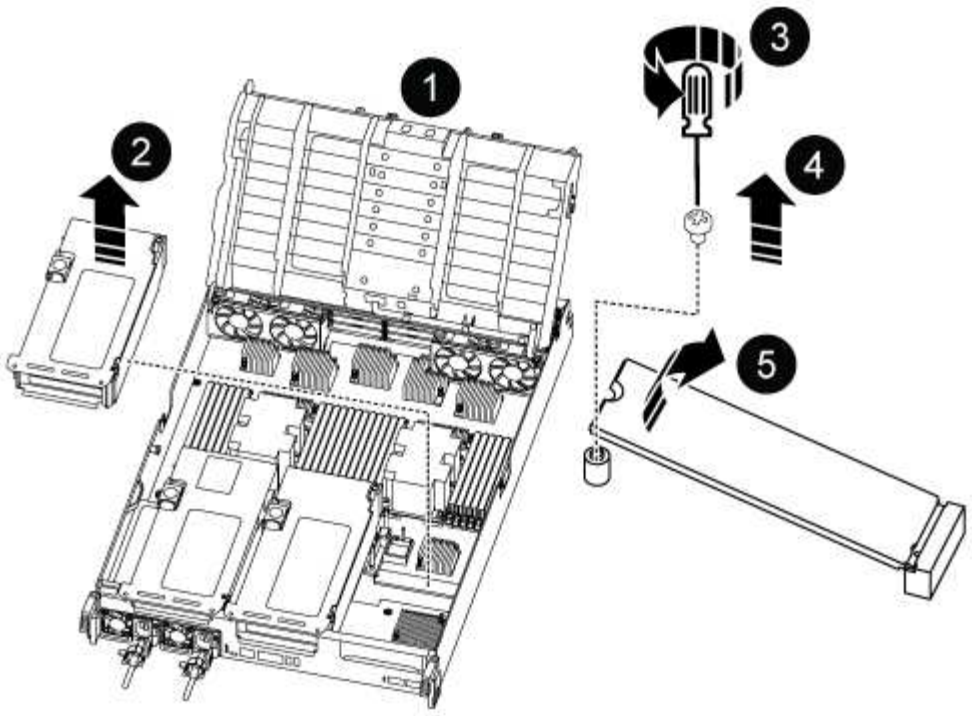

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

2. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.
   b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.
3. Install the replacement boot media into the controller module:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the motherboard.
   c. Secure the boot media to the motherboard using the boot media screw.

      Do not over-tighten the screw or you might damage the boot media.
4. Reinstall the riser into the controller module.

5. Close the air duct:
    a. Rotate the air duct downward.
    b. Slide the air duct toward the risers until it clicks into place.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed is without a boot image so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
    - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
    - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
    a. Download the service image to your work space on your laptop.
    b. Unzip the service image.

        If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

        There are two folders in the unzipped service image file:

        - boot
        - efi

    c. Copy the efi folder to the top directory on the USB flash drive.

        If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

    The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

    a. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct:
    a. Swing the air duct all the way down to the controller module.
    b. Slide the air duct toward the risers until the locking tabs click into place.

c. Inspect the air duct to make sure that it is properly seated and locked into place.

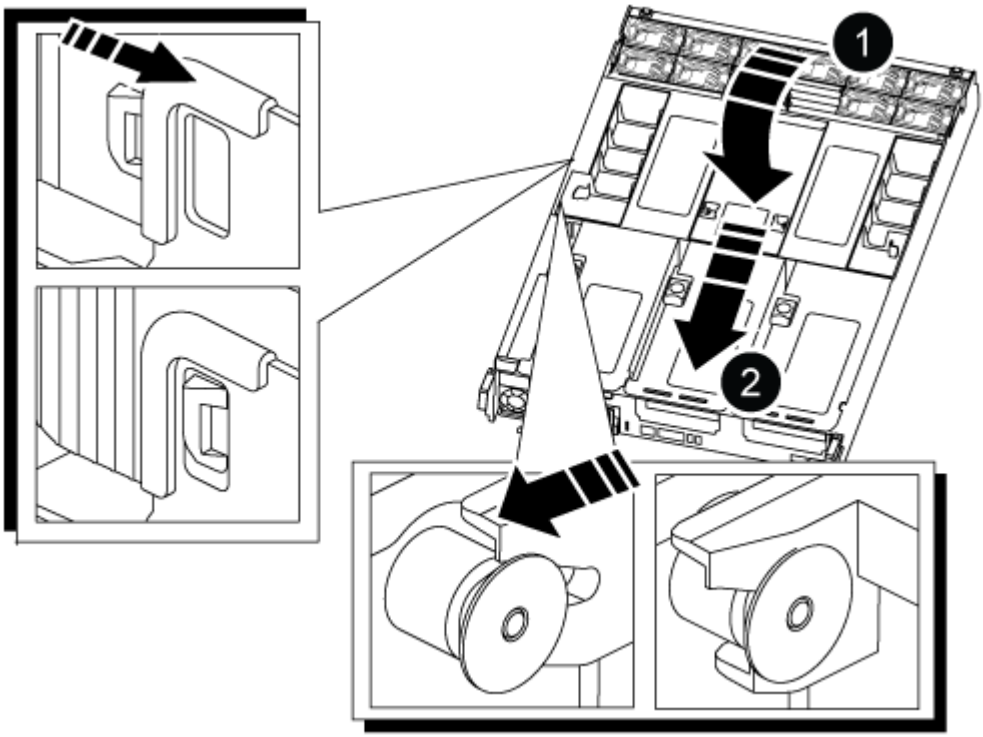

1	Air duct
2	Risers

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.

5. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

6. Gently push the controller module all the way into the system until the controller module locking hooks begin to rise, firmly push on the locking hooks to finish seating the controller module, and then swing the locking hooks into the locked position over the pins on the controller module.
7. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the

controller to boot to LOADER.

Manual boot media recovery from a USB drive - AFF A800

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore encryption - AFF A800

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*******************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*******************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - AFF A800

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Replace the chassis - AFF A800

To replace the chassis, you must move the bezel, controller modules, and NVMe drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in its place.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A800

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - AFF A800

Move the power supplies, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

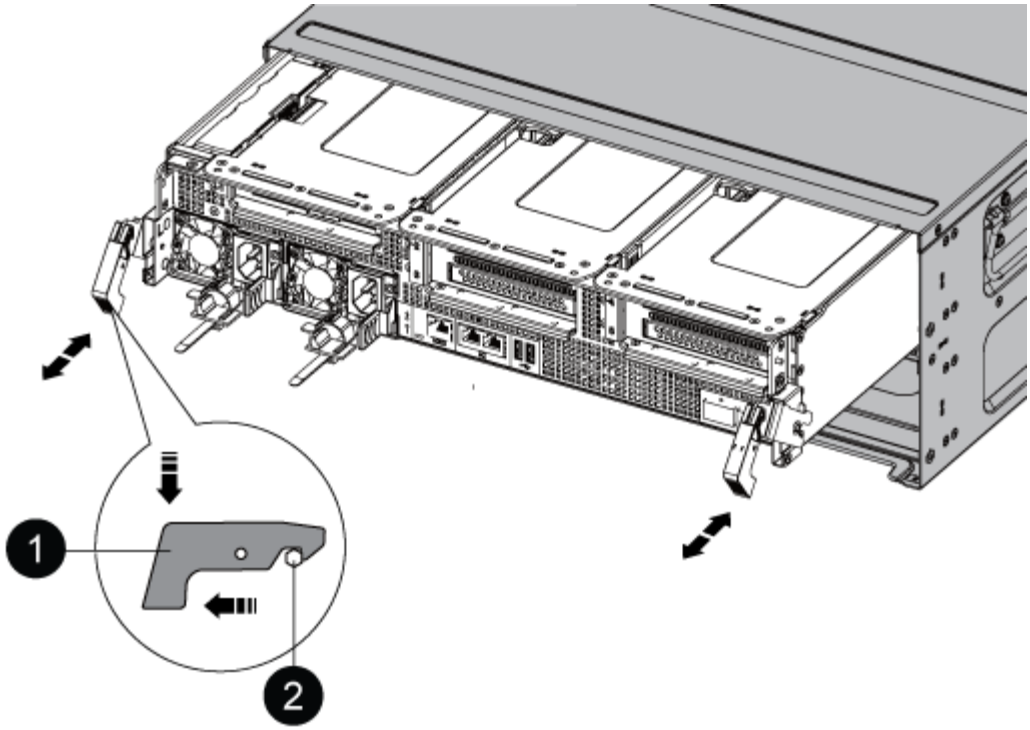

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.

3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.

5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process by pressing `Ctrl-C`.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - AFF A800

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - AFF A800

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Do not downgrade the BIOS version of the *replacement* controller to match the partner controller or the old controller module.

Shut down the impaired controller - AFF A800

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   `cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - AFF A800

To replace the controller, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

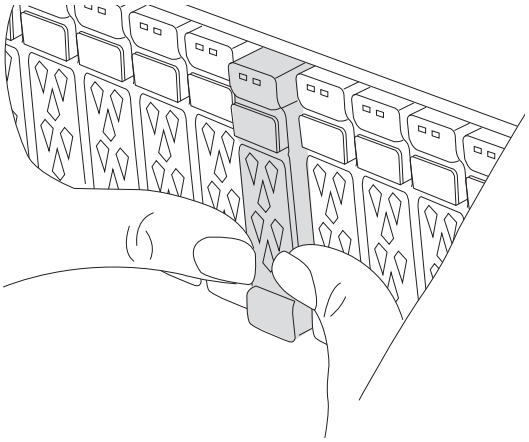

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

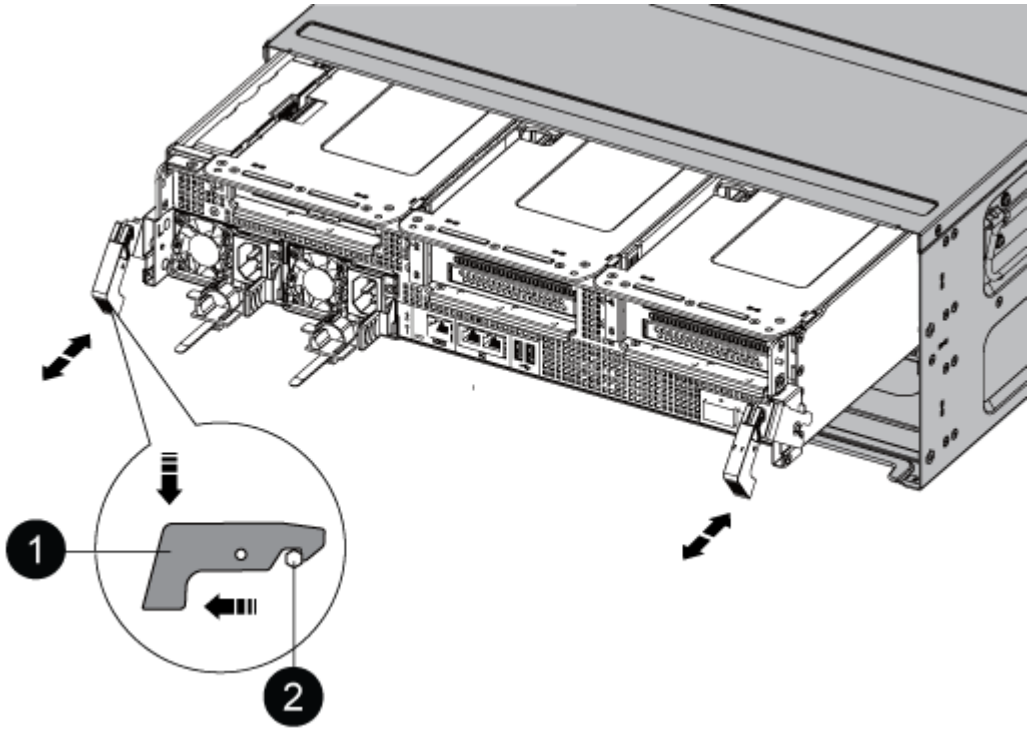

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
   .

9. Place the controller module on a stable, flat surface.
10. Open the controller module air duct:
    a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
    b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

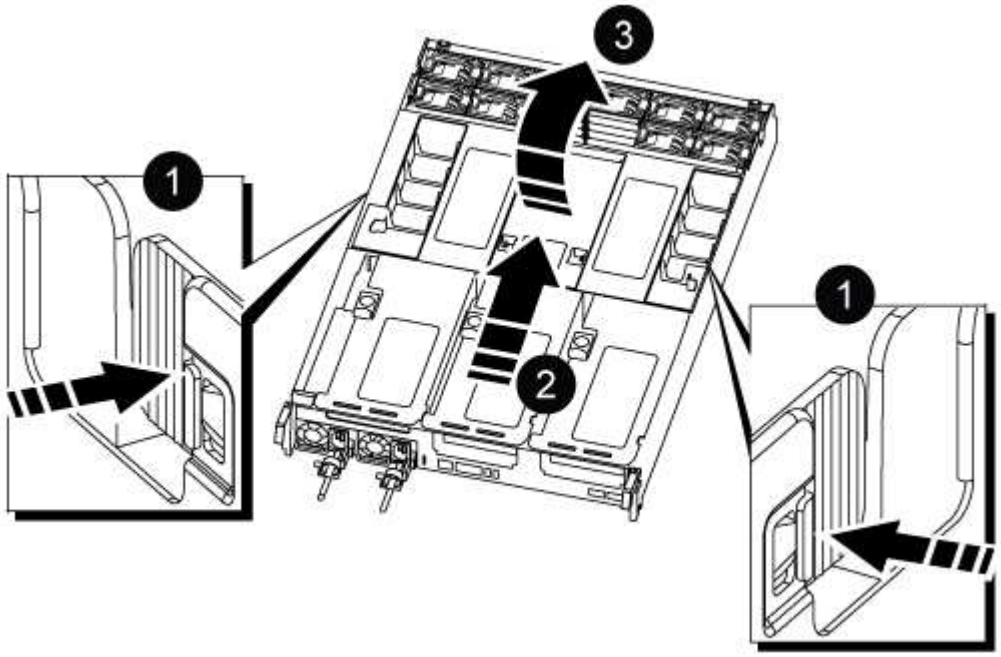

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 2: Move the power supplies

You must move the power supplies from the impaired controller module to the replacement controller module when you replace a controller module.

1. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

> The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

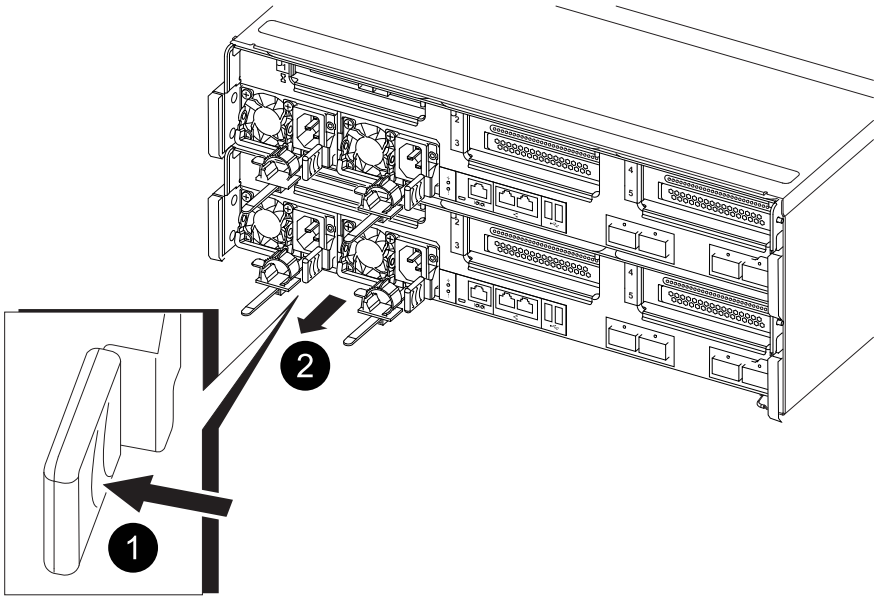

1	Blue power supply locking tab
2	Power supply

2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

> To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

You must move the fans from the impaired controller module to the replacement module when replacing a failed controller module.

1. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

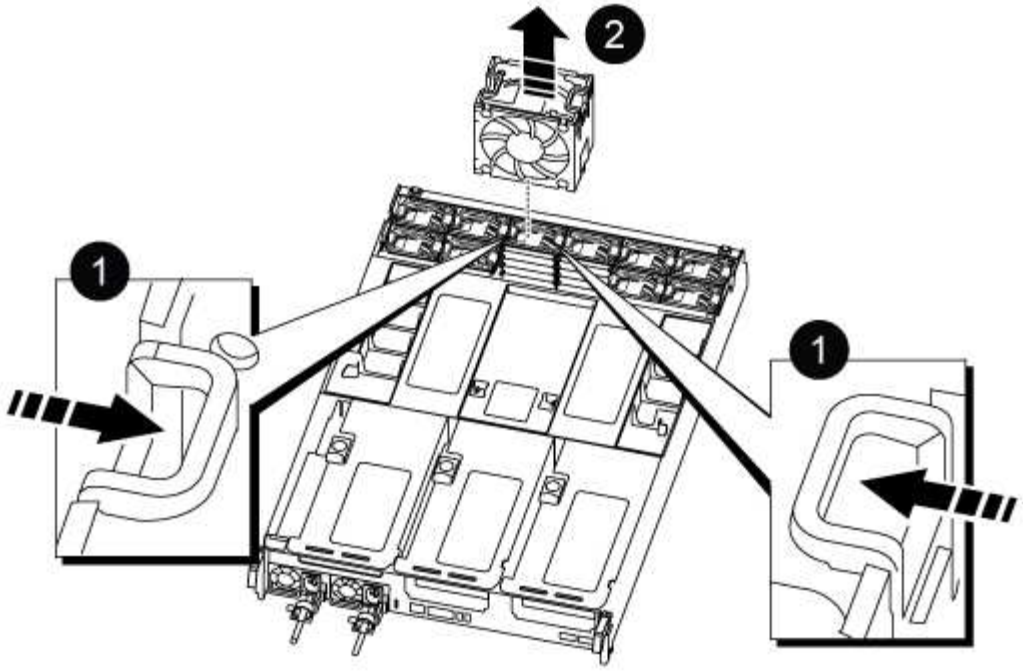

1	Fan locking tabs
2	Fan module

2. Move the fan module to the replacement controller module, and then install the fan module by aligning its edges with the opening in the controller module, and then sliding the fan module into the controller module until the locking latches click into place.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the NVDIMM battery

When replacing the controller module, you must move the NVRAM battery from the impaired controller module to the replacement controller module

1. Open the air duct cover and locate the NVDIMM battery in the riser.

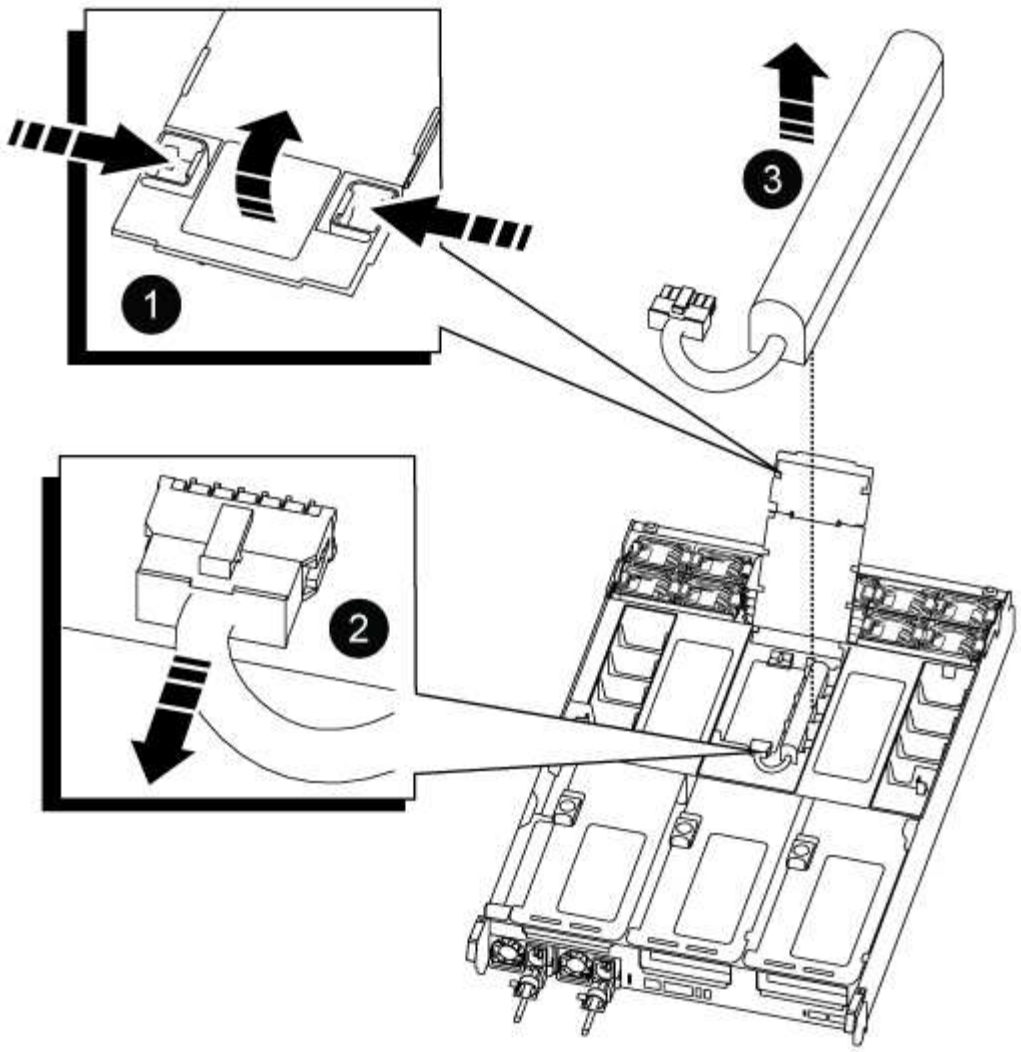

1	Air duct riser
2	NVDIMM battery plug
3	NVDIMM battery pack

Attention: The NVDIMM battery control board LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
3. Grasp the battery and lift the battery out of the air duct and controller module.
4. Move the battery pack to the replacement controller module and then install it in the NVDIMM air duct:
   a. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   b. Plug the battery plug into the riser socket and make sure that the plug locks into place.

Step 5: Remove the PCIe risers

As part of the controller replacement process, you must remove the PCIe modules from the impaired controller module. You must install them into the same location in the replacement controller module once the NVDIMMS and DIMMs have moved to the replacement controller module.

1. Remove the PCIe riser from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

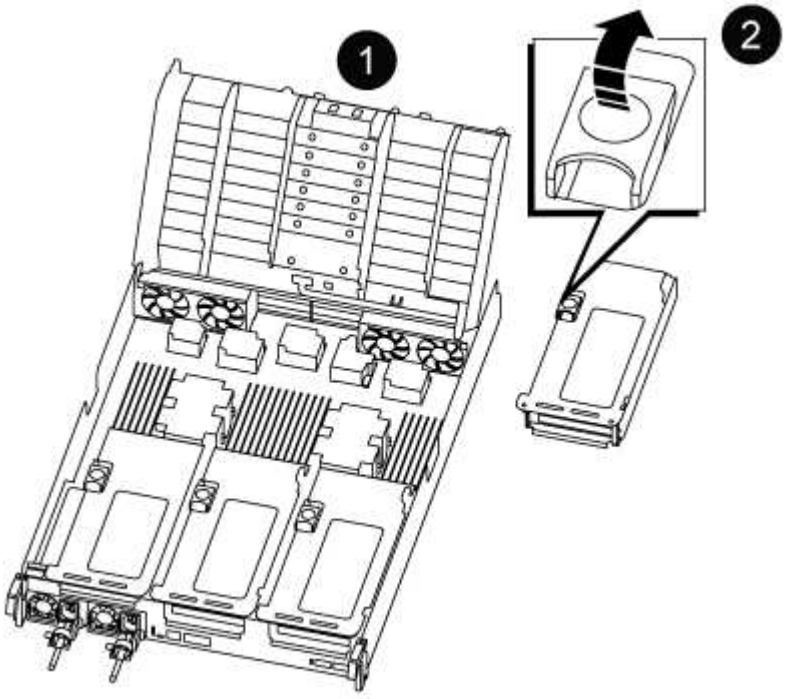

1	Air duct
2	Riser 1 (left riser), Riser 2 (middle riser), and 3 (right riser) locking latches

2. Repeat the preceding step for the remaining risers in the impaired controller module.
3. Repeat the above steps with the empty risers in the replacement controller and put them away.

Step 6: Move system DIMMs

To move the DIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

1. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.

2. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

3. Locate the slot where you are installing the DIMM.
4. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

5. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
6. Repeat these steps for the remaining DIMMs.

Step 7: Move the NVDIMMs

To move the NVDIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

1. Locate the NVDIMMs on your controller module.

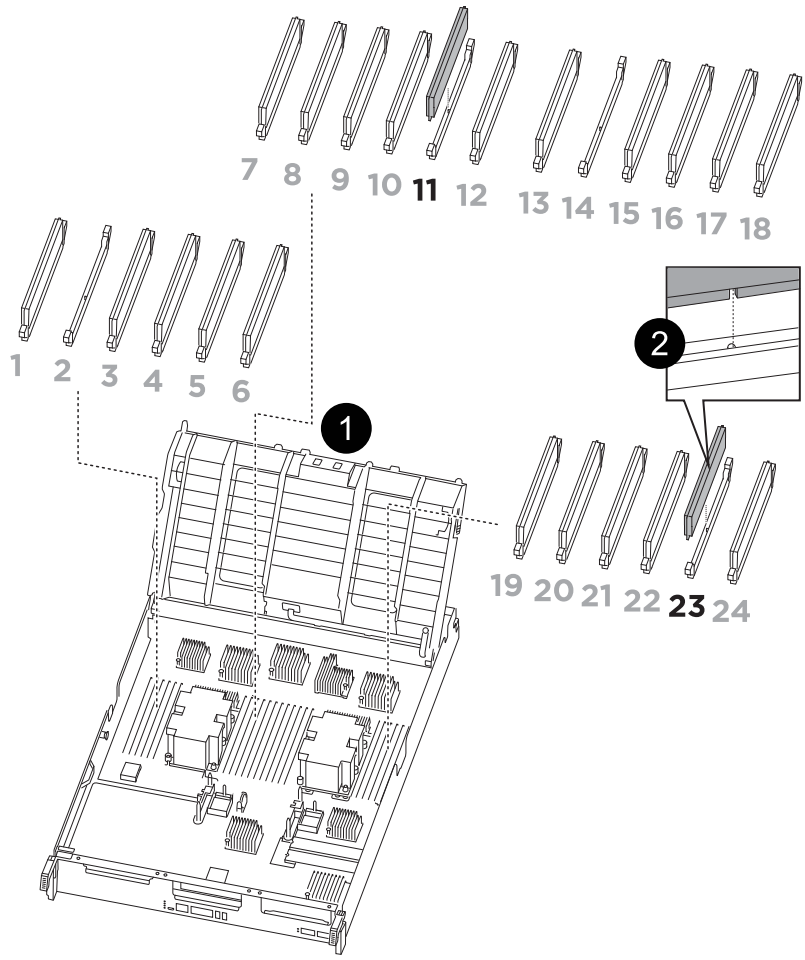

- NVDIMM: SLOTS 11 & 23

1	Air duct
2	NVDIMMs

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

4. Locate the slot where you are installing the NVDIMM.
5. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
7. Repeat the preceding steps to move the other NVDIMM.

Step 8: Move the boot media

You must move the boot media device from the impaired controller and install it in the replacement controller.

The boot media is located under Riser 3.

1. Locate the boot media:

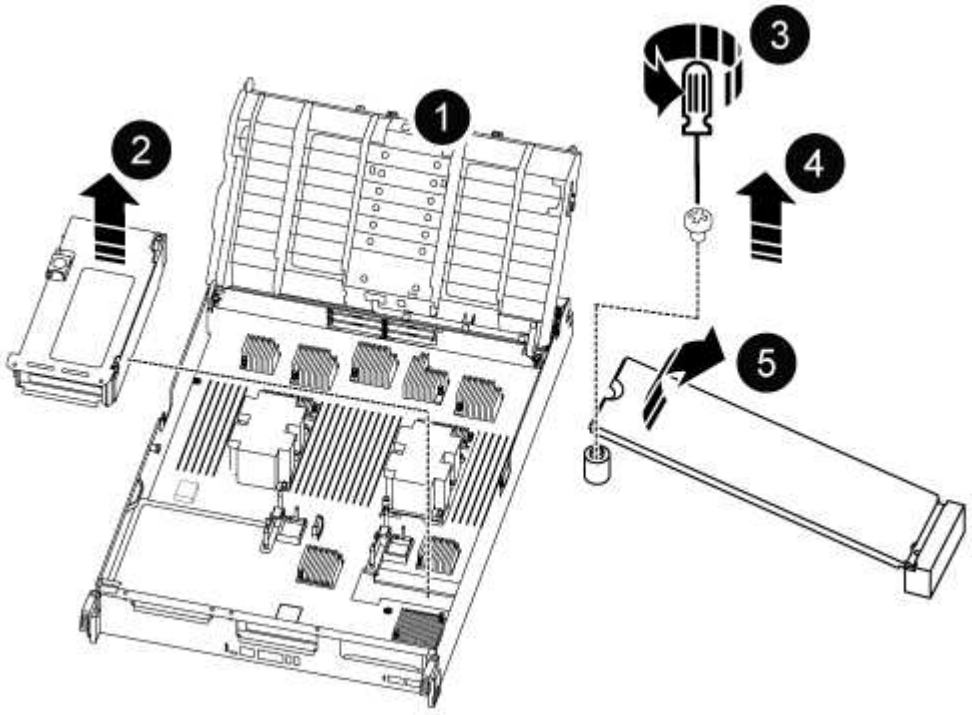

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

2. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.

b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.

3. Move the boot media to the new controller module and install it:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the motherboard.
   c. Secure the boot media to the motherboard using the boot media screw.

      Do not over-tighten the screw or you might damage the boot media.

Step 9: Install the PCIe risers

You install the PCIe risers in the replacement controller module after moving the DIMMs, NVDIMMs, and boot media.

1. Install the riser into the replacement controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

   d. Reinsert any SFP or QSFP modules that were removed from the PCIe cards.
2. Repeat the preceding step for the remaining PCIe risers.

Step 10: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

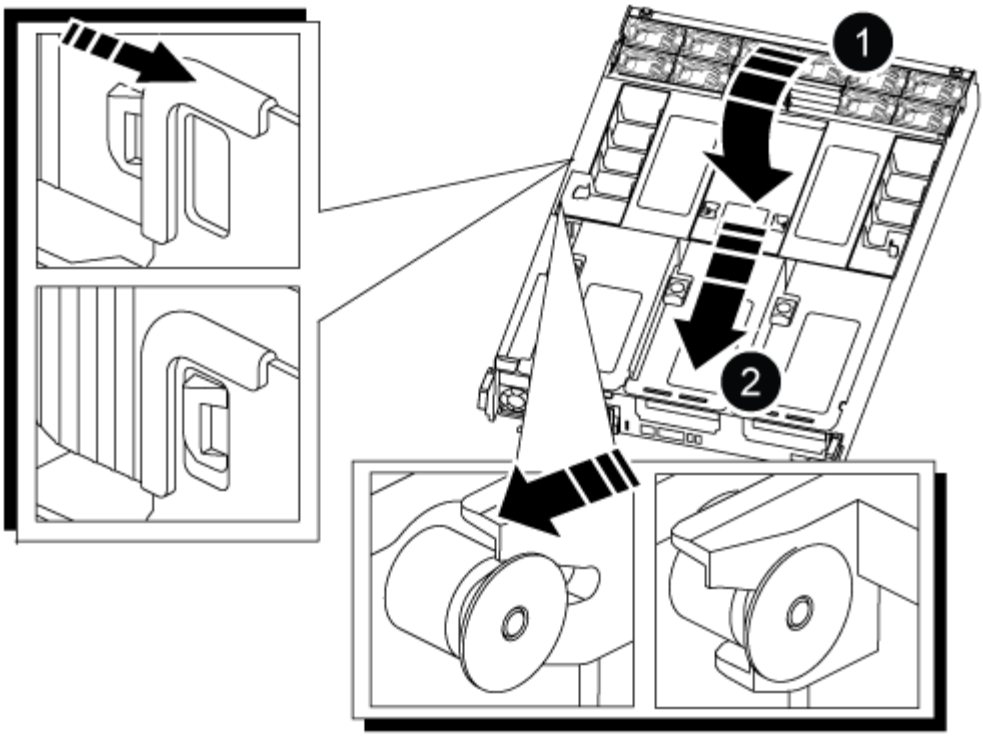

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:

   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

5. Plug the system cables and transceiver modules into the controller module and reinstall the cable management device.

6. Plug the power cables into the power supplies and reinstall the power cable retainers.

   The controller module begins to boot as soon as it is connected to power. Be prepared to interrupt the boot process.

   If your system has DC power supplies, make sure the thumbscrews on the power supply cable are tight.

Restore and verify the system configuration - AFF A800

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the chassis

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA`

state: `ha-config show`

The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF A800

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and

then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF A800

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A800

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

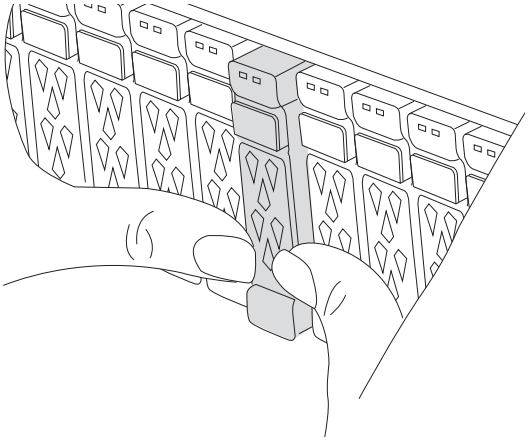

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

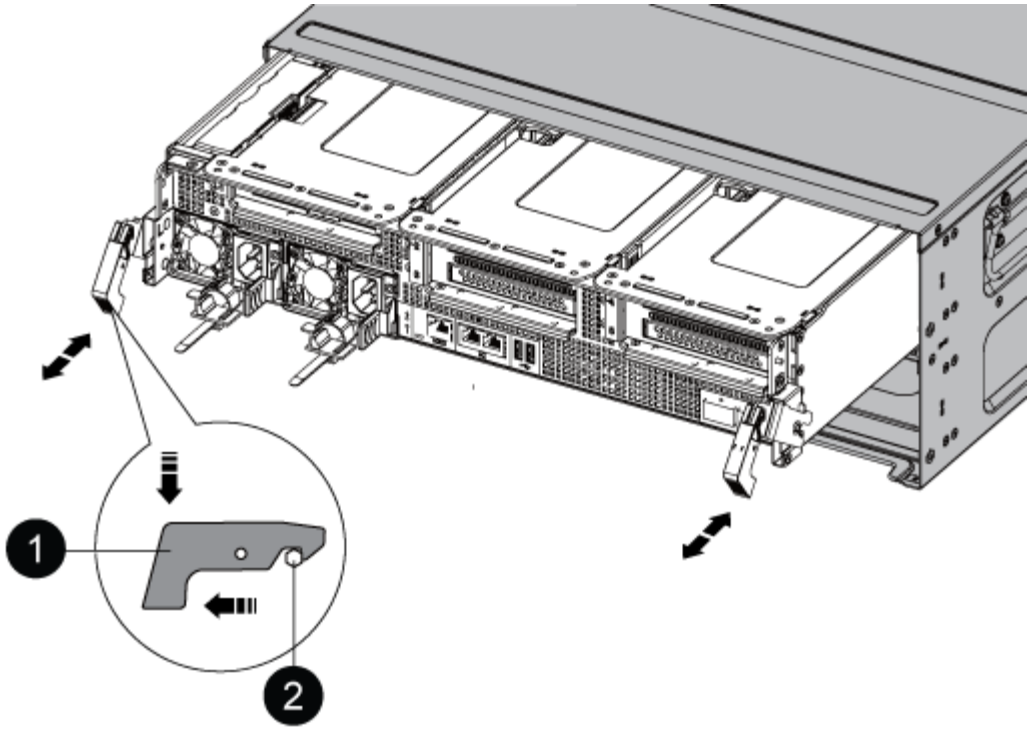

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

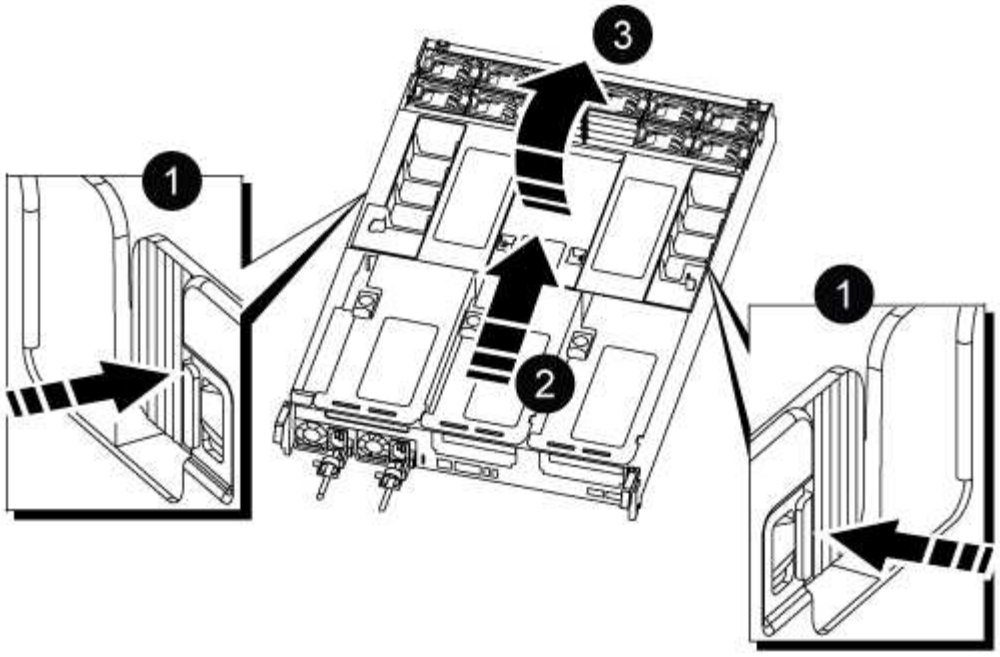

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace a DIMM

To replace a DIMM, you must locate it in the controller module using the DIMM map label on top of the air duct and then replace it following the specific sequence of steps.

1. When removing a DIMM, unlock the locking latch on the applicable riser, and then remove the riser.

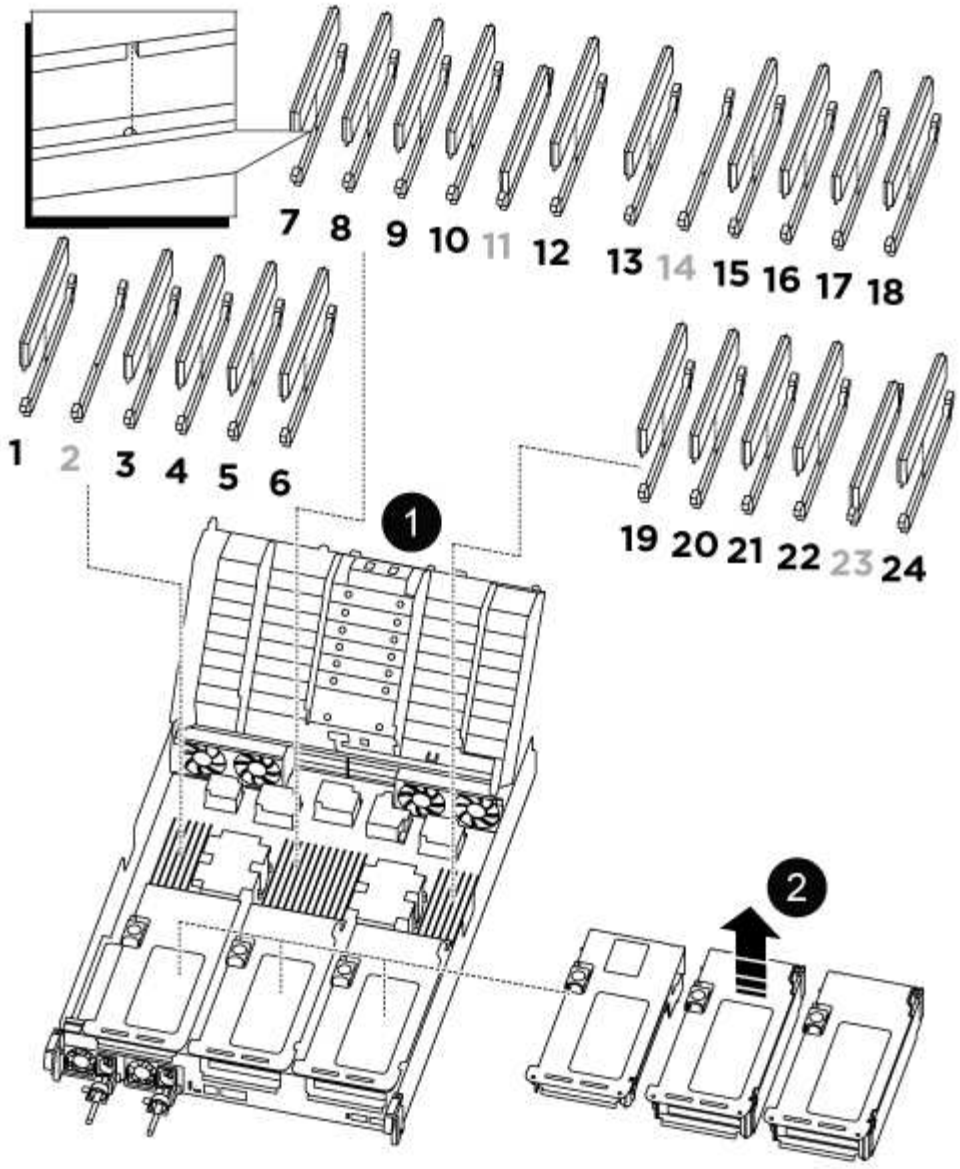

1	Air duct cover
2	Riser 1 and DIMM bank 1, and 3-6
Riser 2 and DIMM bank 7-10, 12-13, and 15-18	Riser 3 and DIMM 19 -22 and 24

Note: Slot 2 and 14 are left empty. Do not attempt to install DIMMs into these slots.

2. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

4. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

5. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Reinstall any risers that you removed from the controller module.
8. Close the air duct.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

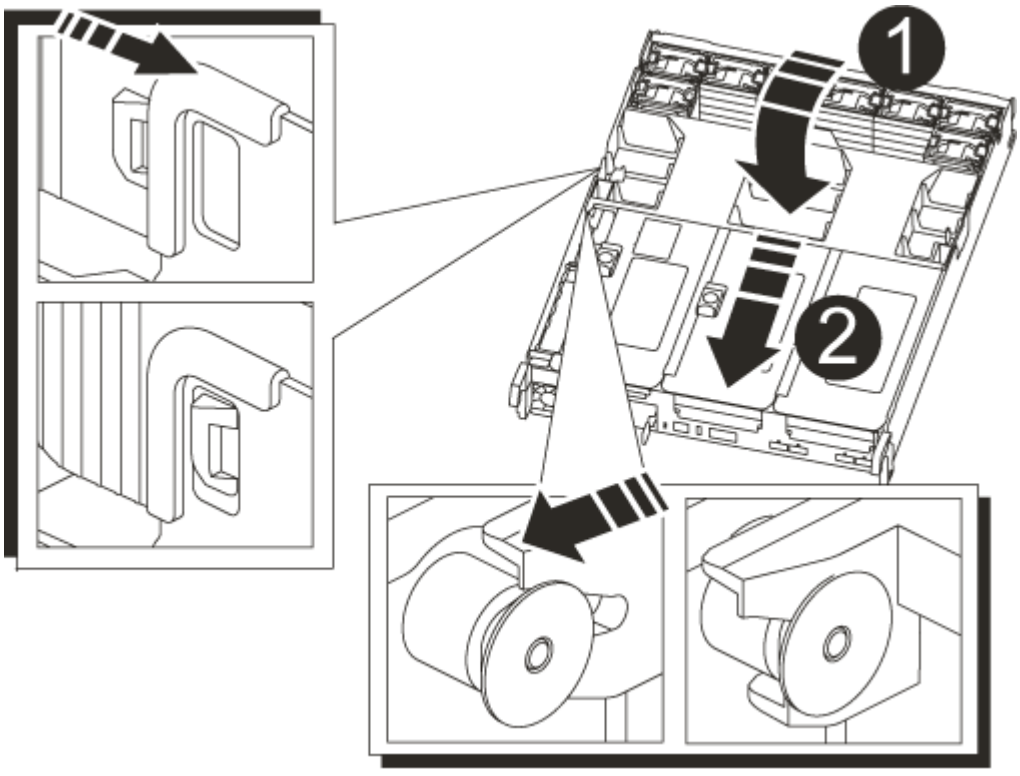

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - AFF A800

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before

replacing a drive.

- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  > Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan - AFF A800

To replace a fan, remove the failed fan module and replace it with a new fan module.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a fan module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

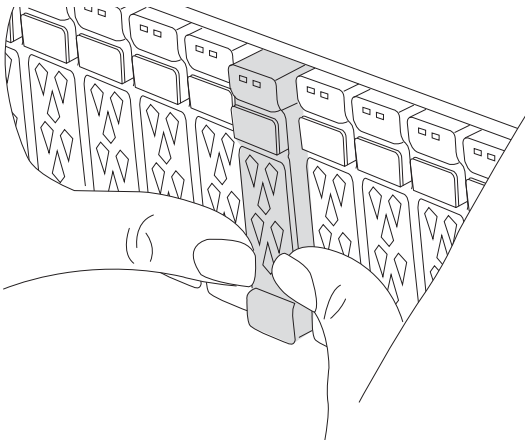

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

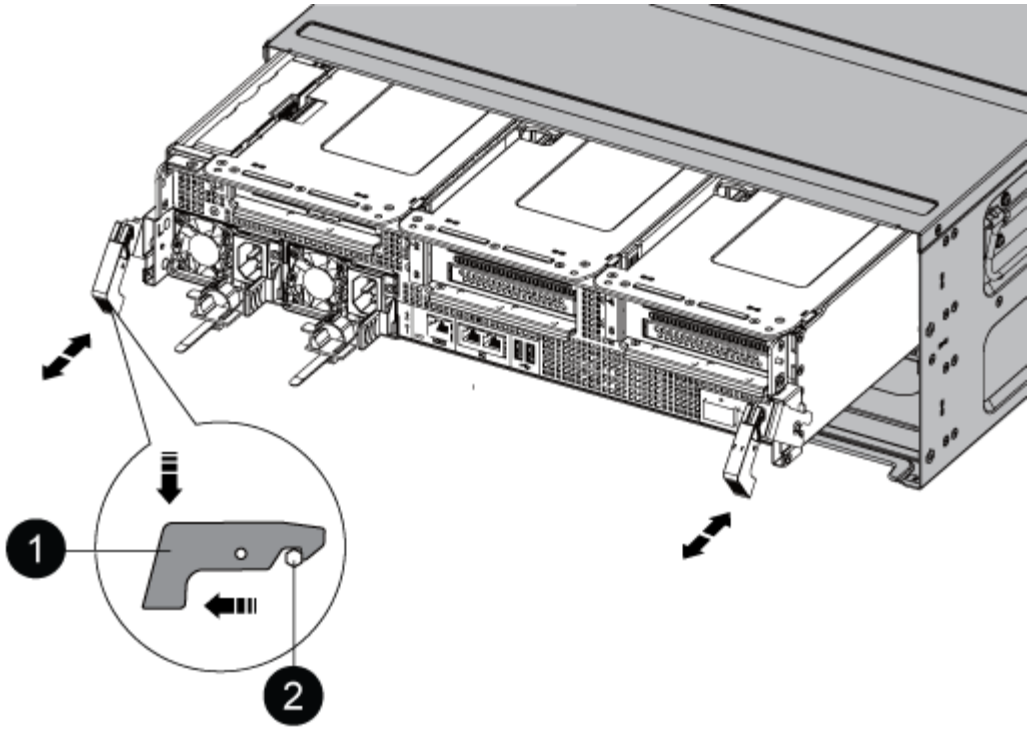

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Set the controller module aside in a safe place.

Step 3: Replace a fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
2. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

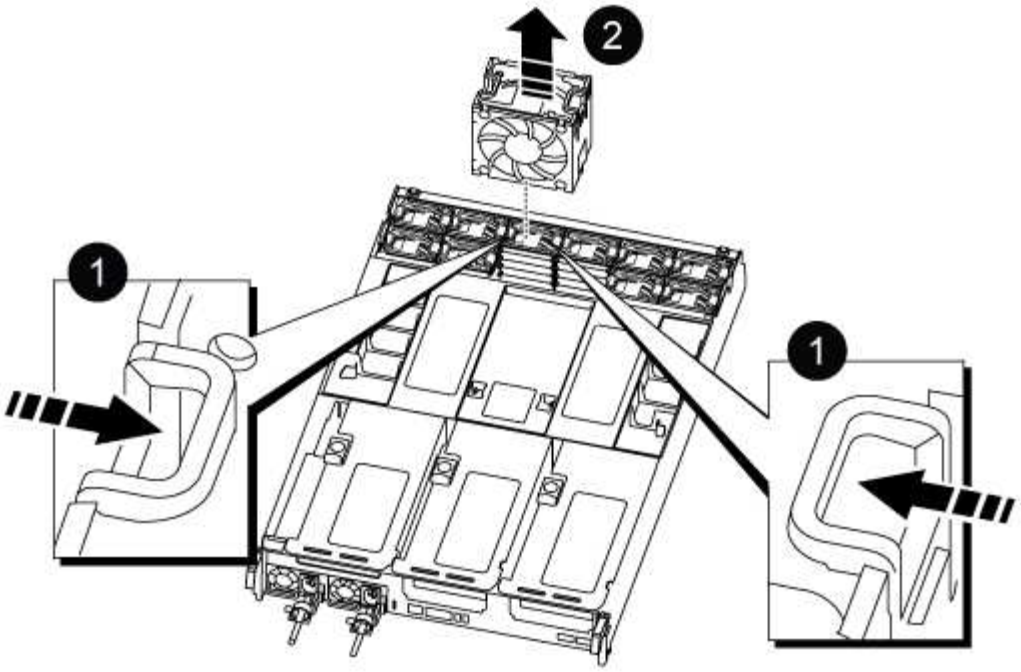

1	Fan locking tabs
2	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module until the locking latches click into place.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.
3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -controller local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an NVDIMM - AFF A800

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

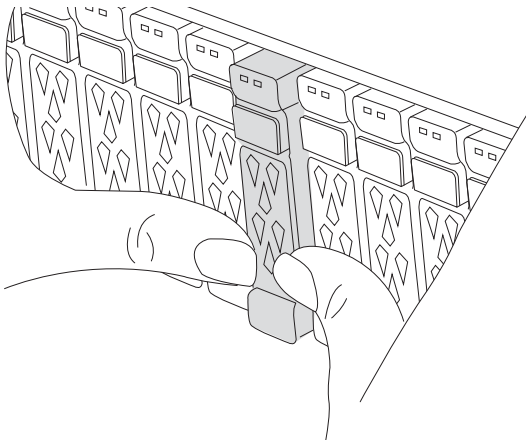

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.

5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

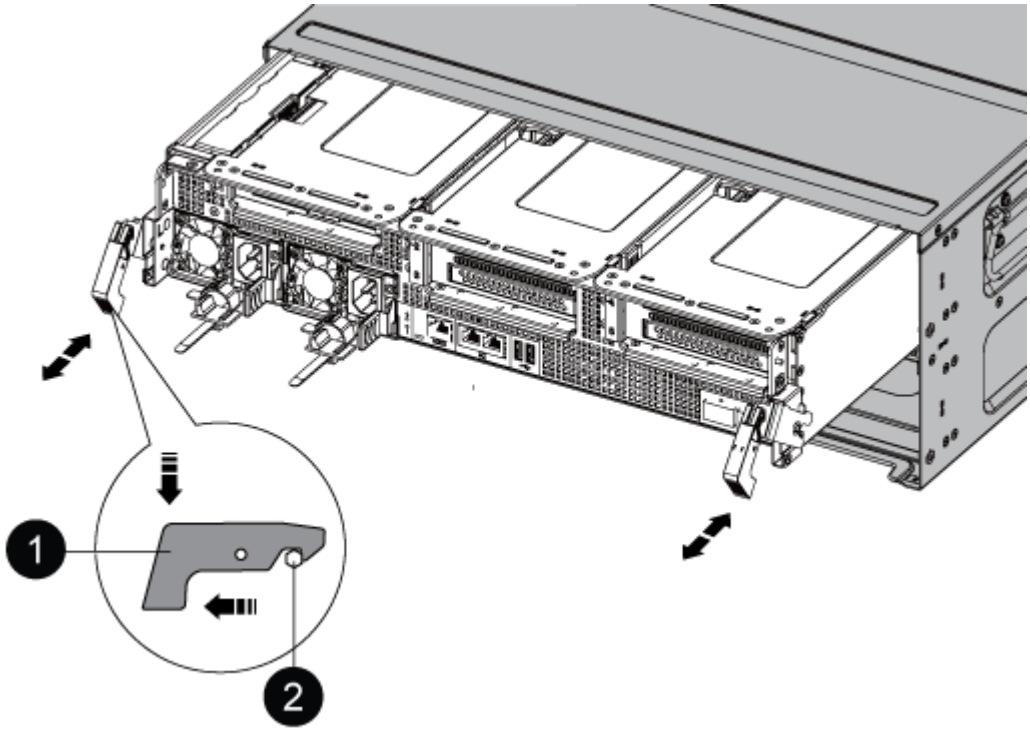

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:

   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

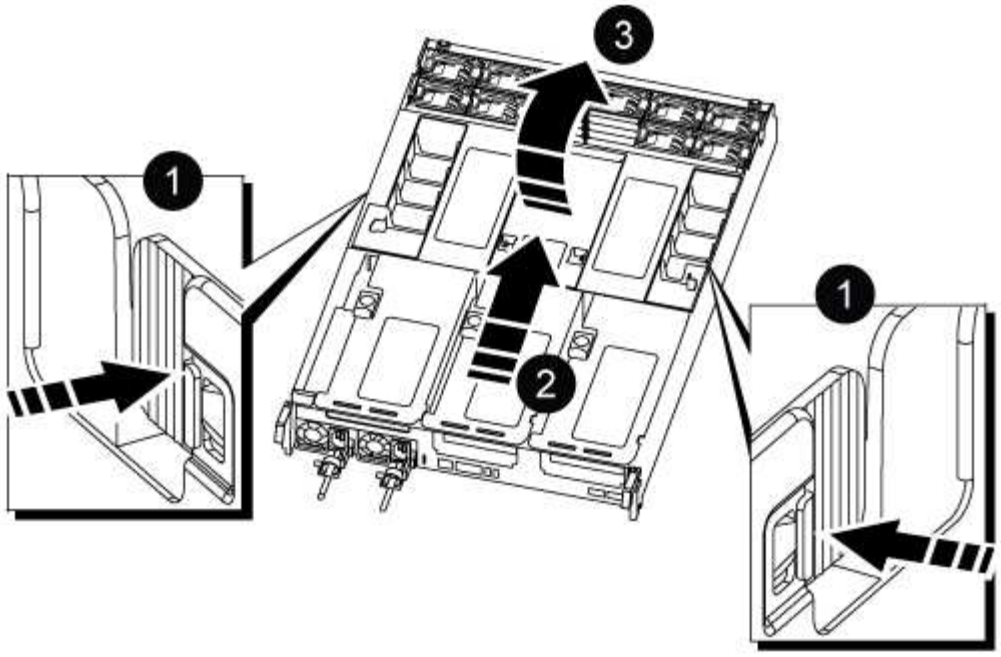

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the NVDIMM map label on top of the air duct, and then replace it following the specific sequence of steps.

1. If you are removing or moving an NVDIMM, unlock the locking latch on the riser, and then remove the applicable riser.

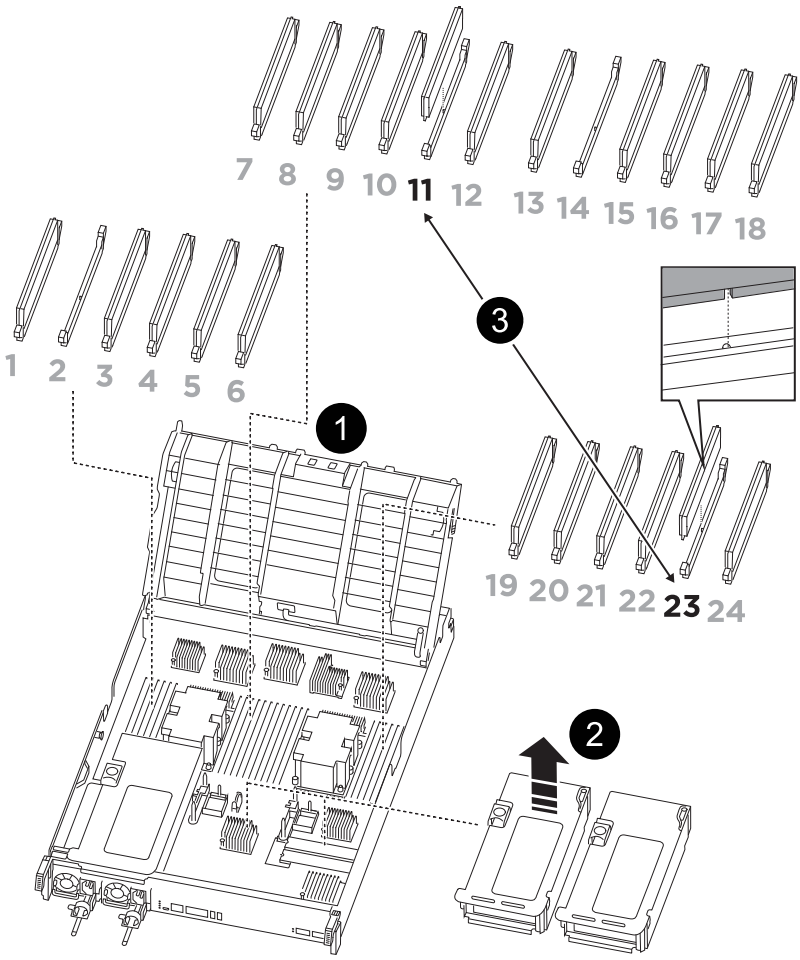

1	Air duct cover
2	Riser 2
3	NVDIMM in slots 11 and 23

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

4. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

5. Locate the slot where you are installing the NVDIMM.

6. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
8. Reinstall any risers that you removed from the controller module.
9. Close the air duct.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

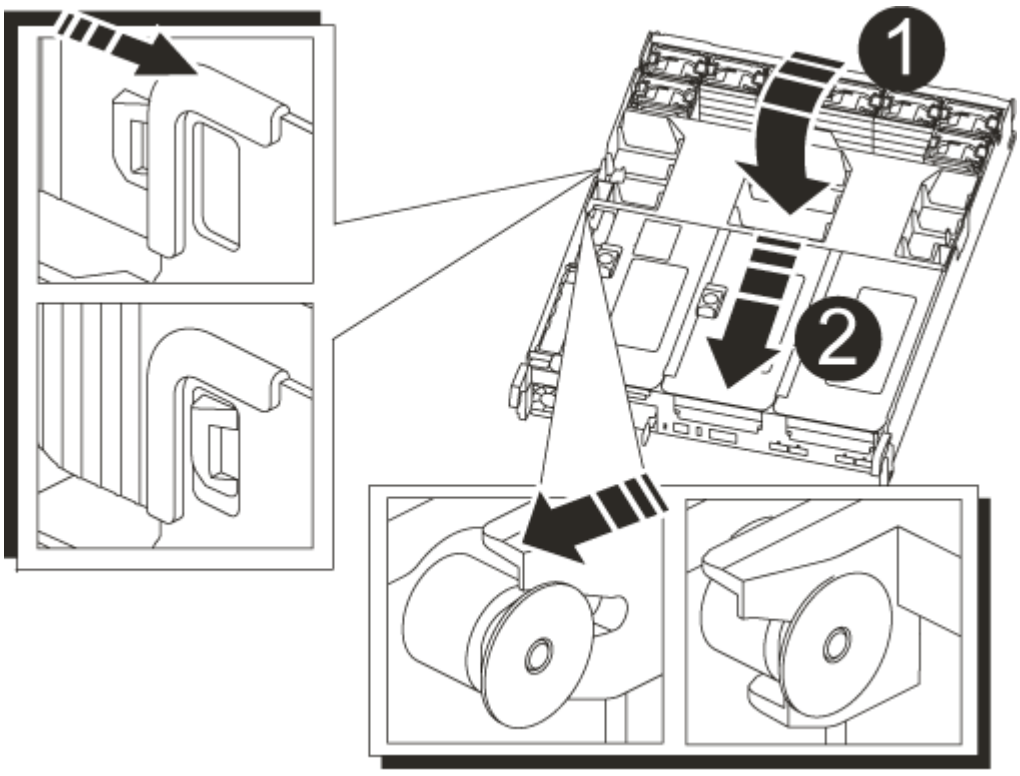

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - AFF A800

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be

resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

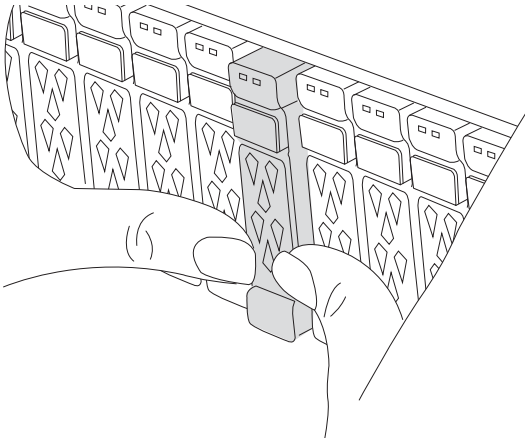

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

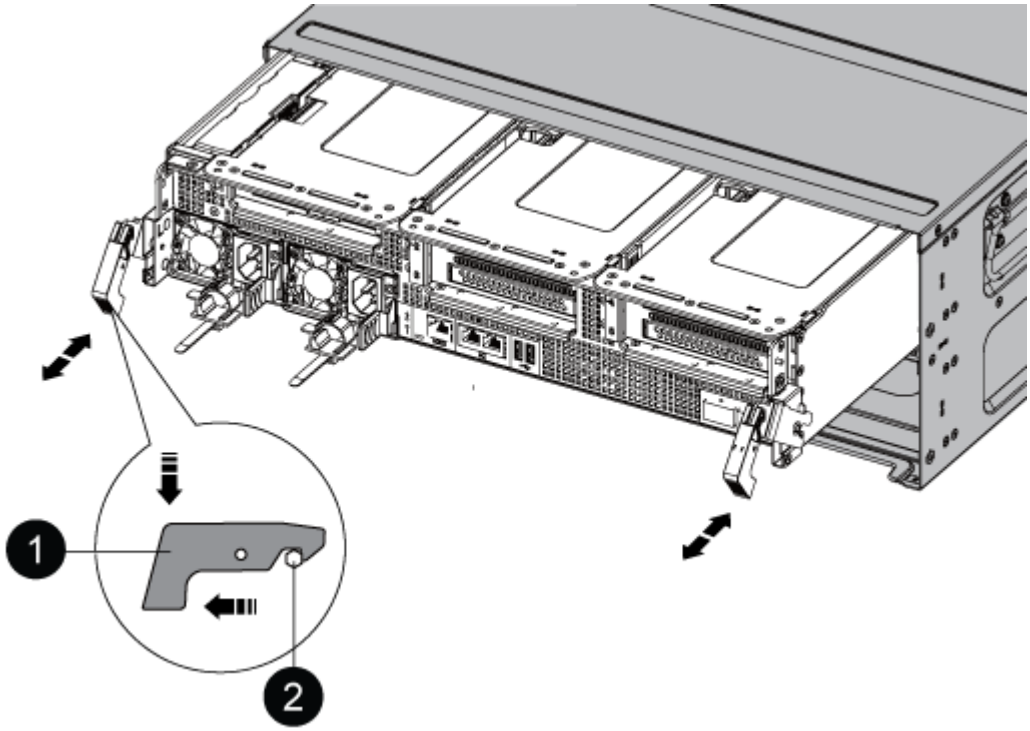

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Set the controller module aside in a safe place.

Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module.

1. Open the air duct cover and locate the NVDIMM battery in the riser.

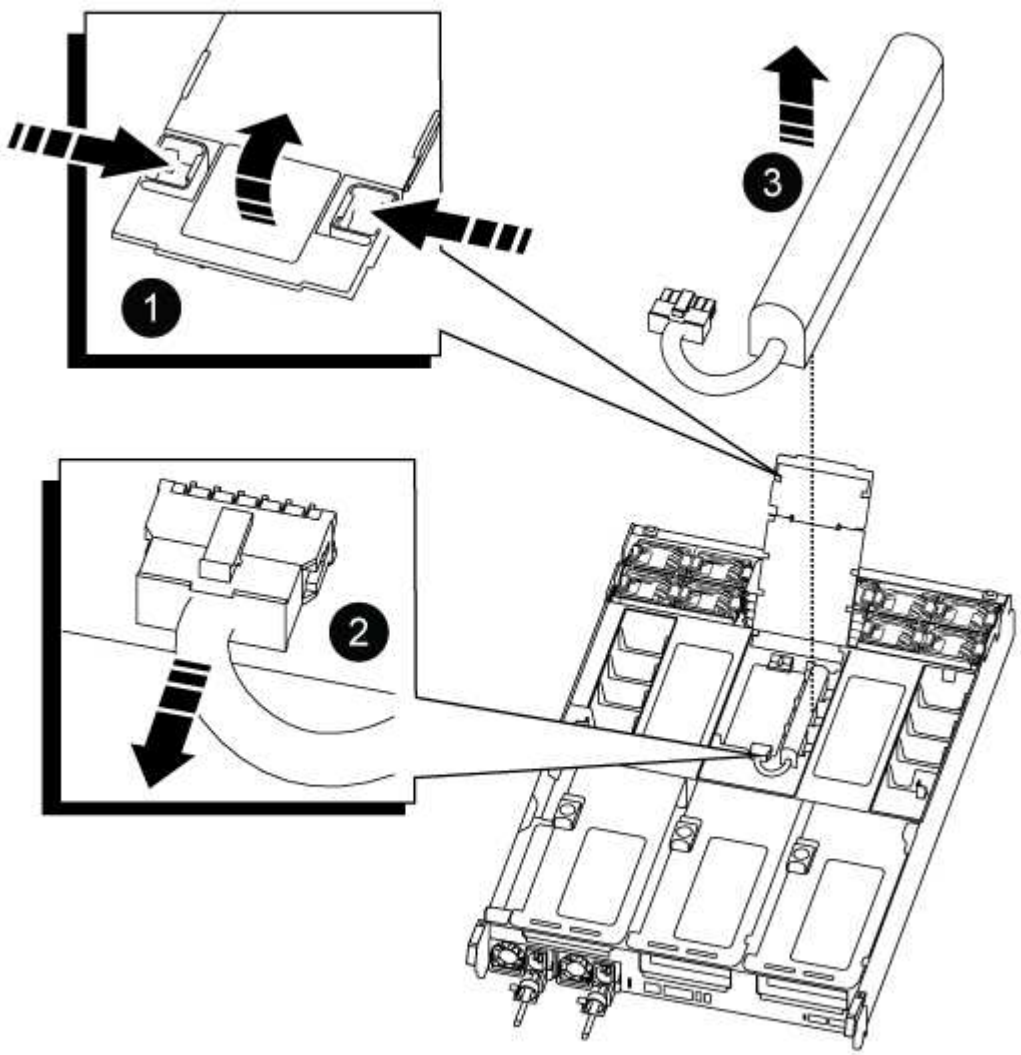

1	Air duct riser

2	NVDIMM battery plug
3	NVDIMM battery pack

Attention: The NVDIMM battery control board LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

1. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
2. Grasp the battery and lift the battery out of the air duct and controller module, and then set it aside.
3. Remove the replacement battery from its package.
4. Install the replacement battery pack in the NVDIMM air duct:
   a. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   b. Plug the battery plug into the riser socket and make sure that the plug locks into place.
5. Close the NVDIMM air duct.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe card - AFF A800

To replace a PCIe card, you must disconnect the cables from the cards, remove the SFP and QSFP modules from the cards before removing the riser, reinstall the riser, and then reinstall the SFP and QSFP modules before cabling the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

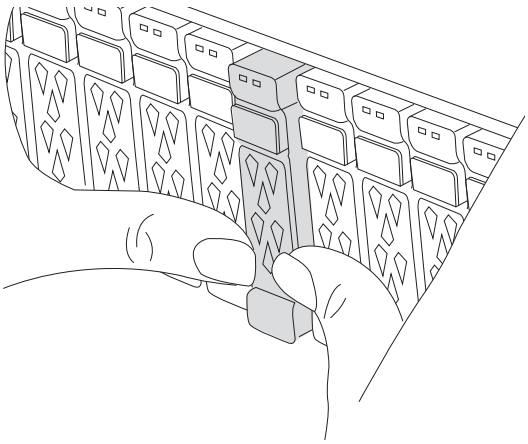

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

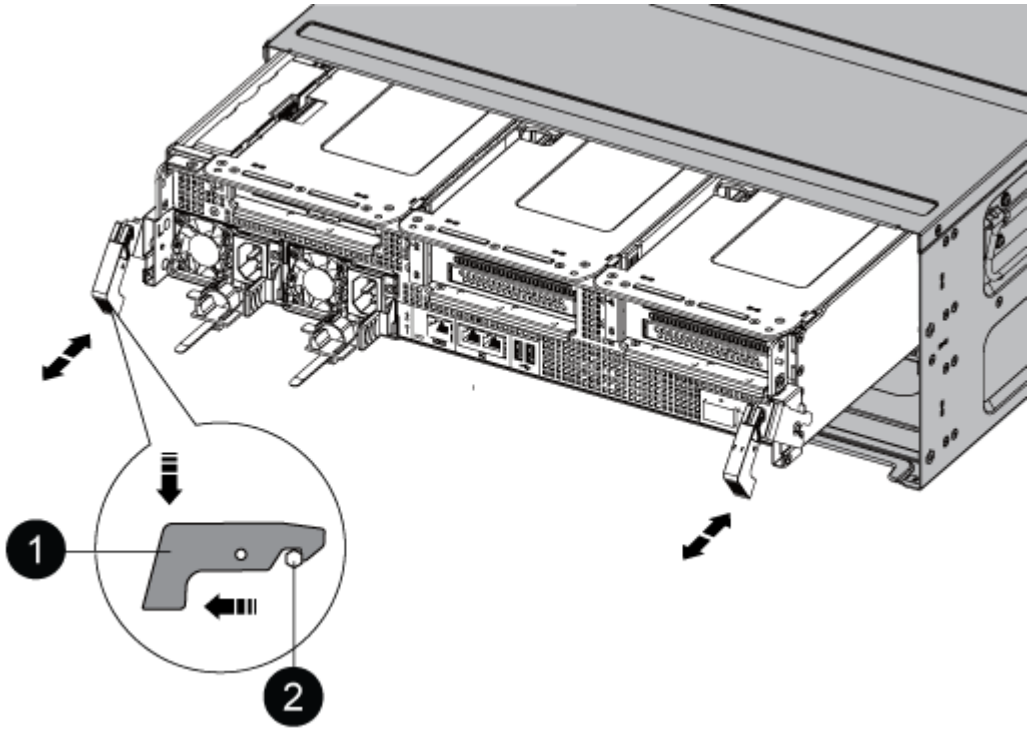

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

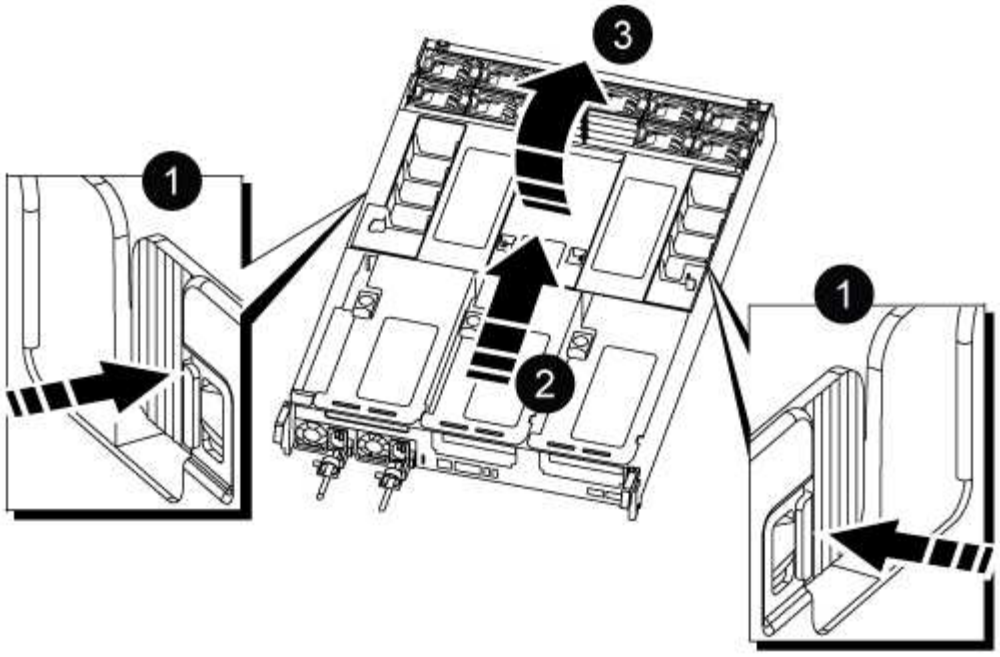

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace a PCIe card

To replace a PCIe card, you must remove the cabling and any QSFPs and SFPs from the ports on the PCIe cards in the target riser, remove the riser from the controller module, remove and replace the PCIe card, reinstall the riser and any QSFPs and SFPs onto the ports, and cable the ports.

1. Determine if the card you are replacing is from Riser 1 or if it is from Riser 2 or 3.
   - If you are replacing the 100GbE PCIe card in Riser 1, use Steps 2 - 3 and Steps 6 - 7.
   - If you are replacing a PCIe card from Riser 2 or 3, use Steps 4 through 7.
2. Remove Riser 1 from the controller module:
   a. Remove the QSFP modules that might be in the PCIe card.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

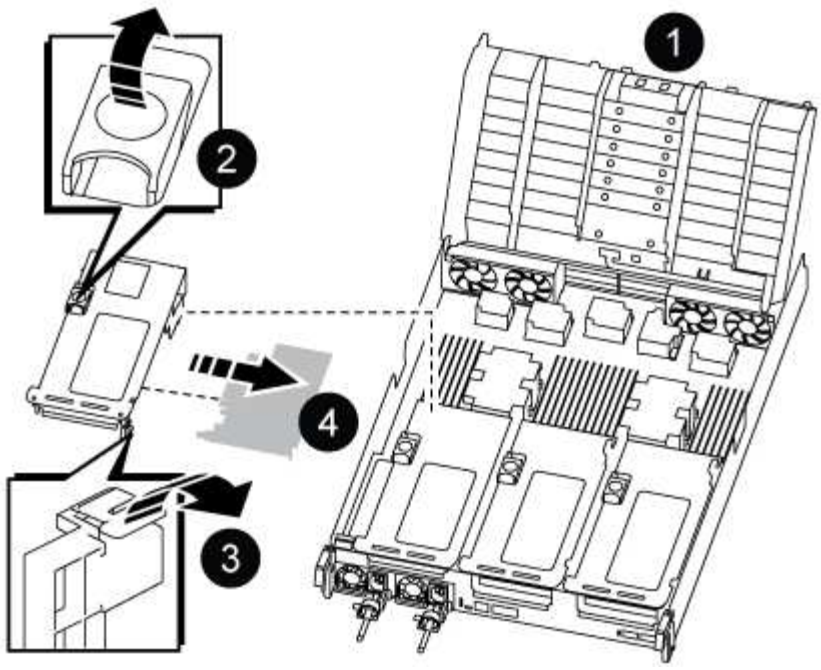

1	Air duct
2	Riser locking latch
3	Card locking bracket
4	Riser 1 (left riser) with 100GbE PCIe card in slot 1.

3. Remove the PCIe card from Riser 1:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. Remove the PCIe card from the riser.
4. Remove the PCIe riser from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

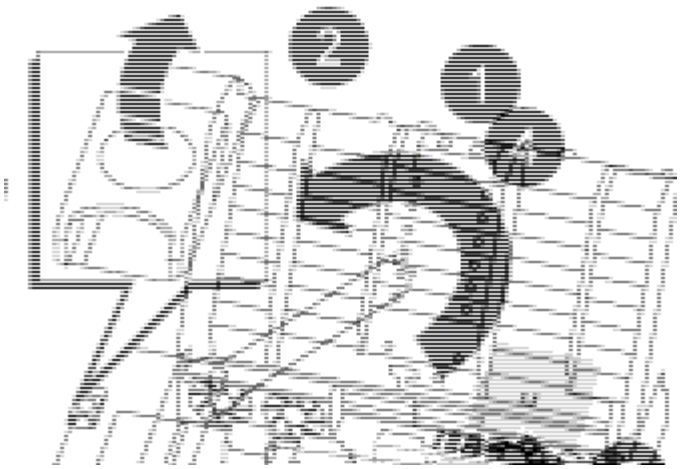

1	Air duct
2	Riser 2 (middle riser) or 3 (right riser) locking latch
3	Card locking bracket
4	Side panel on riser 2 or 3
5	PCIe cards in riser 2 or 3

5. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe cards.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. Swing the side panel off the riser.
   d. Remove the PCIe card from the riser.
6. Install the PCIe card into the same slot in the riser:
   a. Align the card with the card socket in the riser, and then slide it squarely into the socket in the riser.

      Make sure that the card is completely and squarely seated into the riser socket.

   b. For Riser 2 or 3, close the side panel.
   c. Swing the locking latch into place until it clicks into the locked position.
7. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the

controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

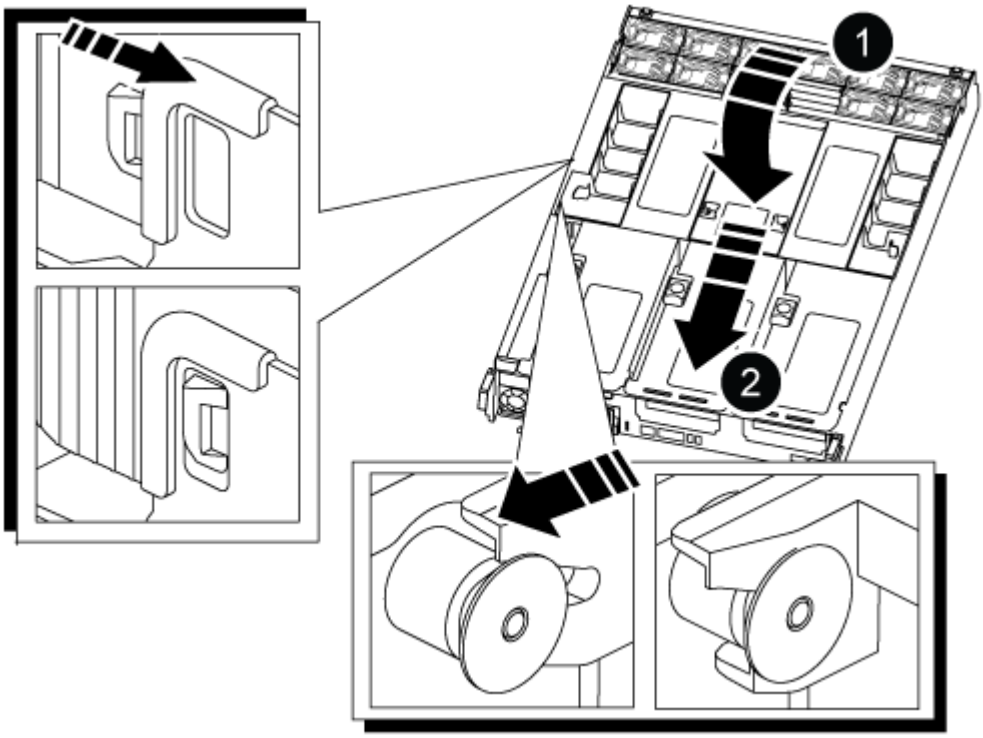

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
5. Plug the system cables and transceiver modules into the controller module and reinstall the cable management device.
6. Plug the power cables into the power supplies and reinstall the power cable retainers.

   The controller module begins to boot as soon as it is connected to power. Be prepared to interrupt the boot process.

   > If your system has DC power supplies, make sure the thumbscrews on the power supply cable are tight.

7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - AFF A800

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

About this task

This procedure is written for replacing one PSU at a time.

> It is a best practice to replace the PSU within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded PSU until the PSU is replaced.

> Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
   b. Unplug the power cable from the power source.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

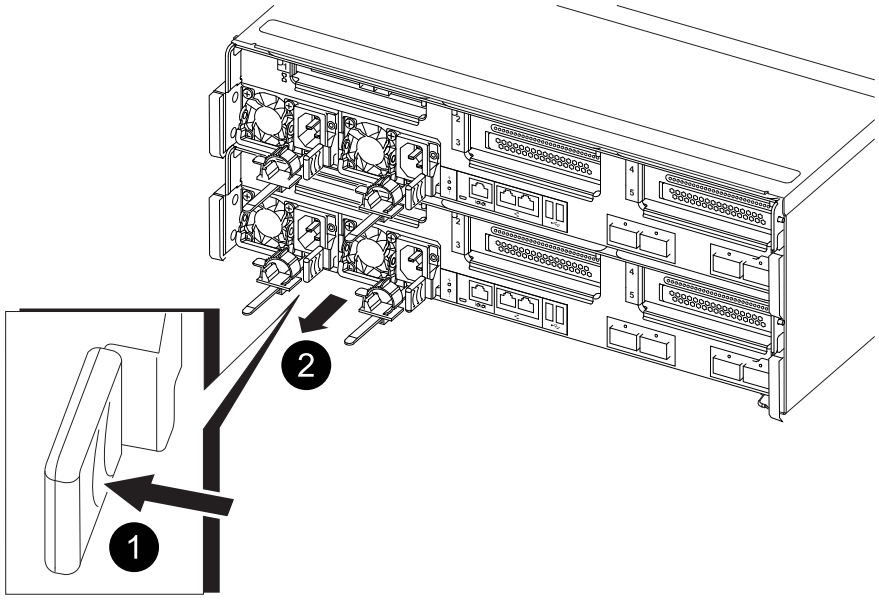

1	Blue PSU locking tab
2	Power supply

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU and the power source.
   b. Secure the power cable to the PSU using the power cable retainer.

   Once power is restored to the PSU, the status LED should be green.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

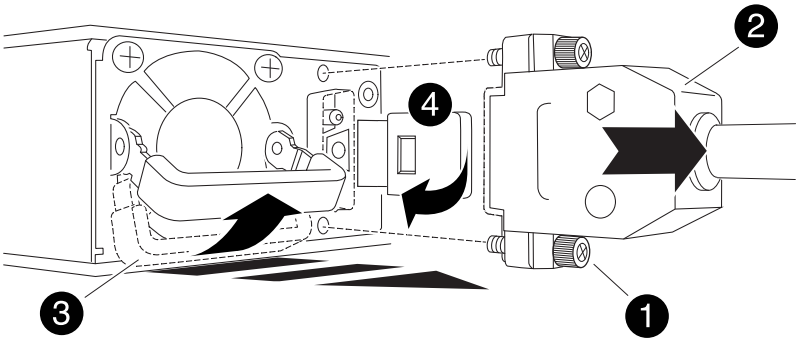

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle

4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A800

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

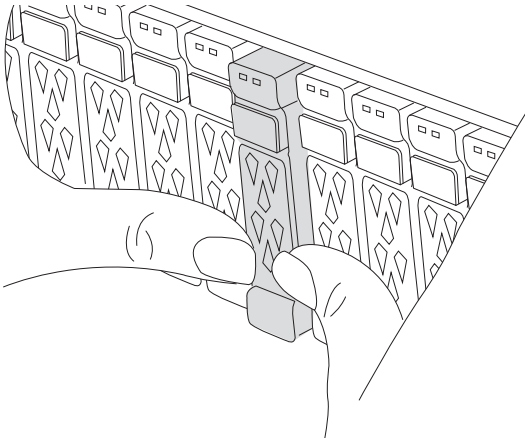

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

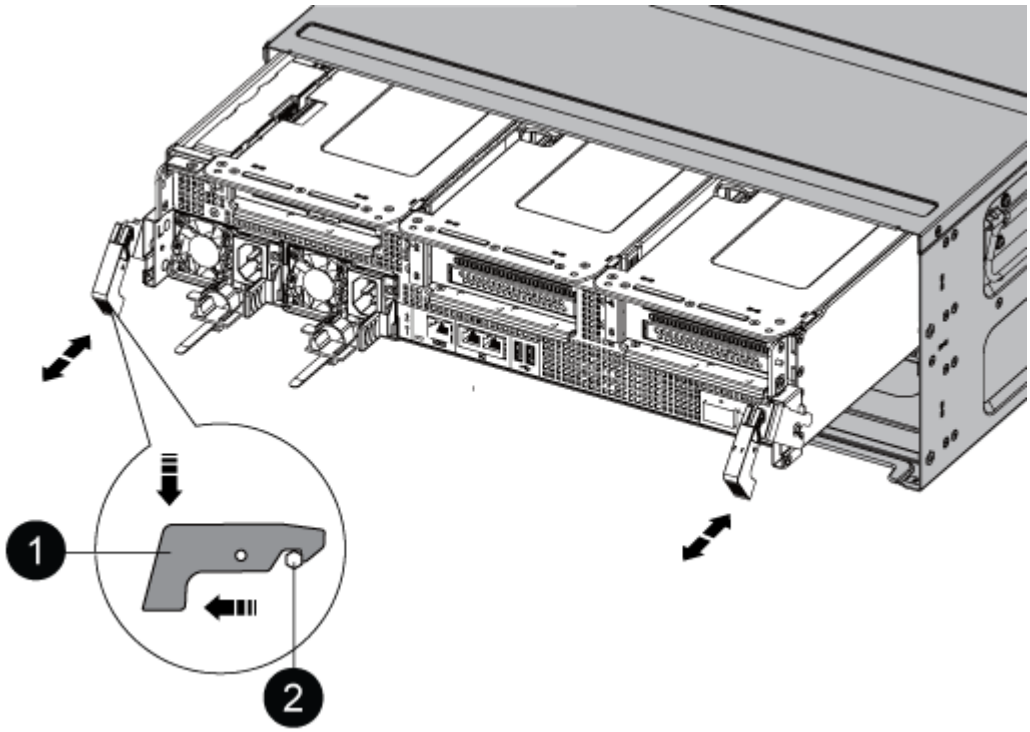

1	Locking latch
2	Locking pin

1. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

2. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

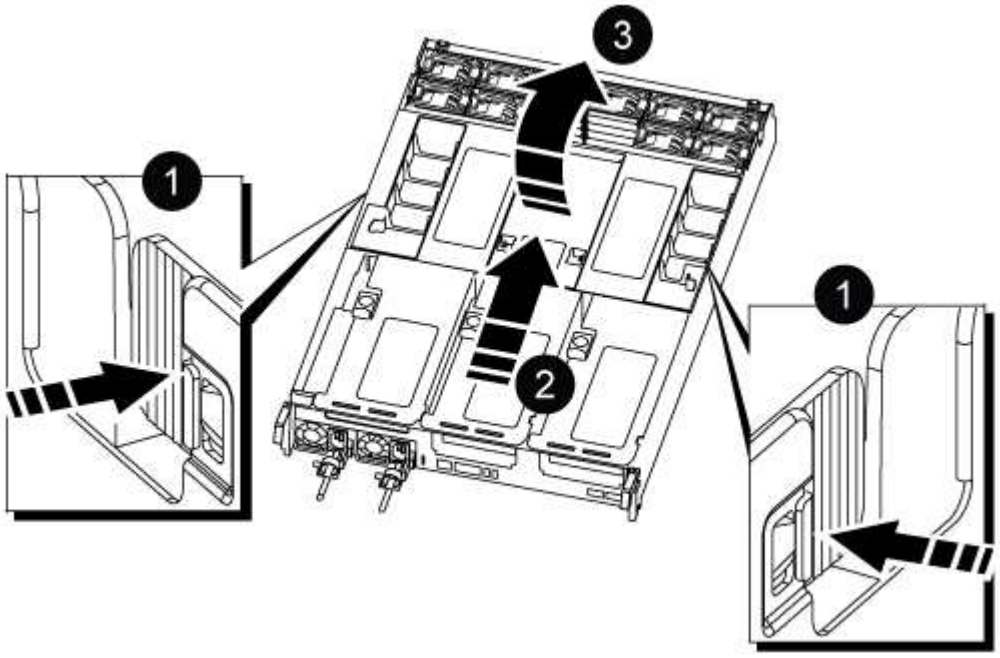

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace the RTC battery

Original controller

1. Remove PCIe riser 2 (middle riser) from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

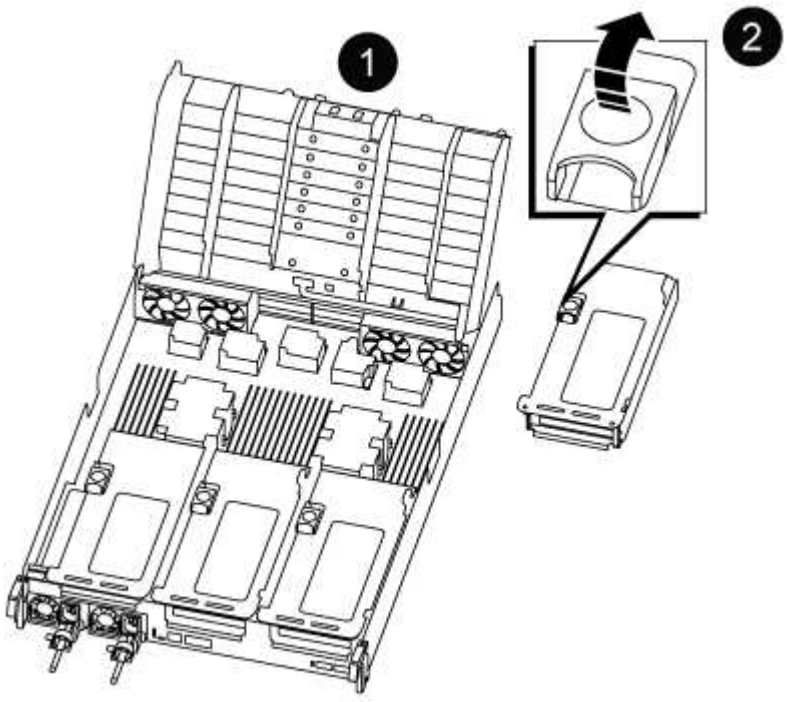

1	Air duct
2	Riser 2 (middle riser) locking latch

2. Locate the RTC battery under Riser 2.

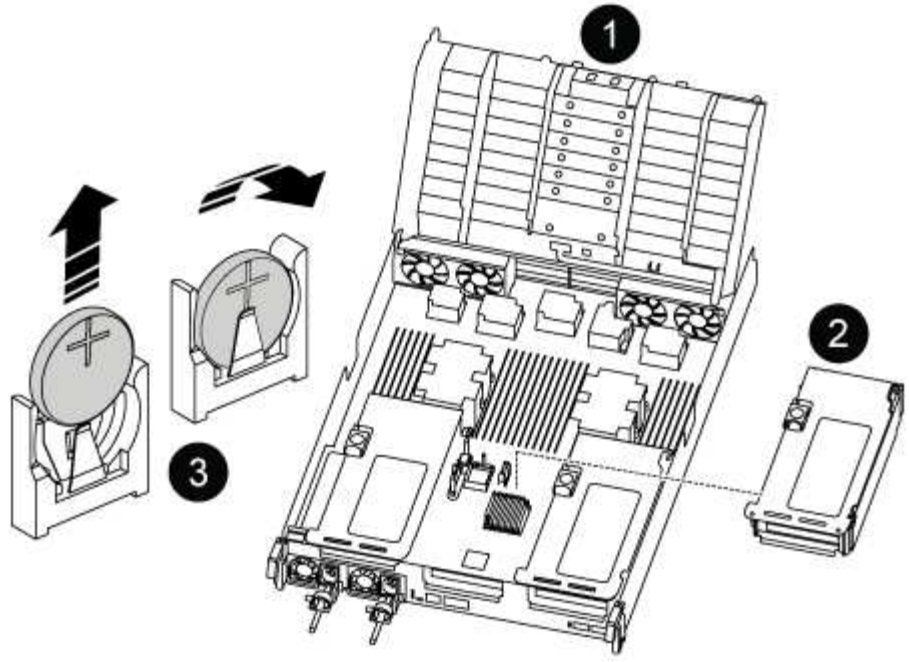

1	Air duct
2	Riser 2
3	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
7. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

VER2 controller

1. Locate the RTC battery near the DIMMs.

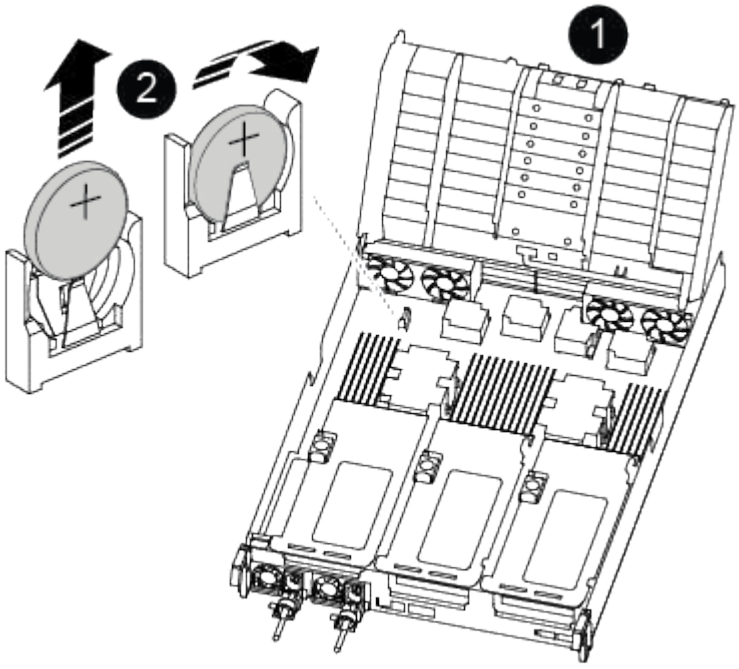

1	Air duct
2	RTC battery and housing

2. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

3. Remove the replacement battery from the antistatic shipping bag.
4. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
5. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Halt the controller at the LOADER prompt.
5. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF A900 systems

Install and setup

Maintain

Maintain AFF A900 hardware

Maintain the hardware of your AFF A900 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the AFF A900 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF A900 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the storage system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
DCPM	The DCPM (destage controller power module) contains the NVRAM11 battery.
Fan	The fan cools the controller.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.

LED USB	The LED USB module provides connectivity to console ports and system status.
NVRAM	The NVRAM module (Non-Volatile Random Access Memory) allows the controller to retain data across power cycles or system reboots, while the NVRAM DIMM maintains NVRAM settings.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - AFF A900

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your AFF A900 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - AFF A900

Before replacing the boot media in your AFF A900, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - AFF A900

Shut down the impaired controller in your AFF A900 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - AFF A900

The boot media in your AFF A900 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

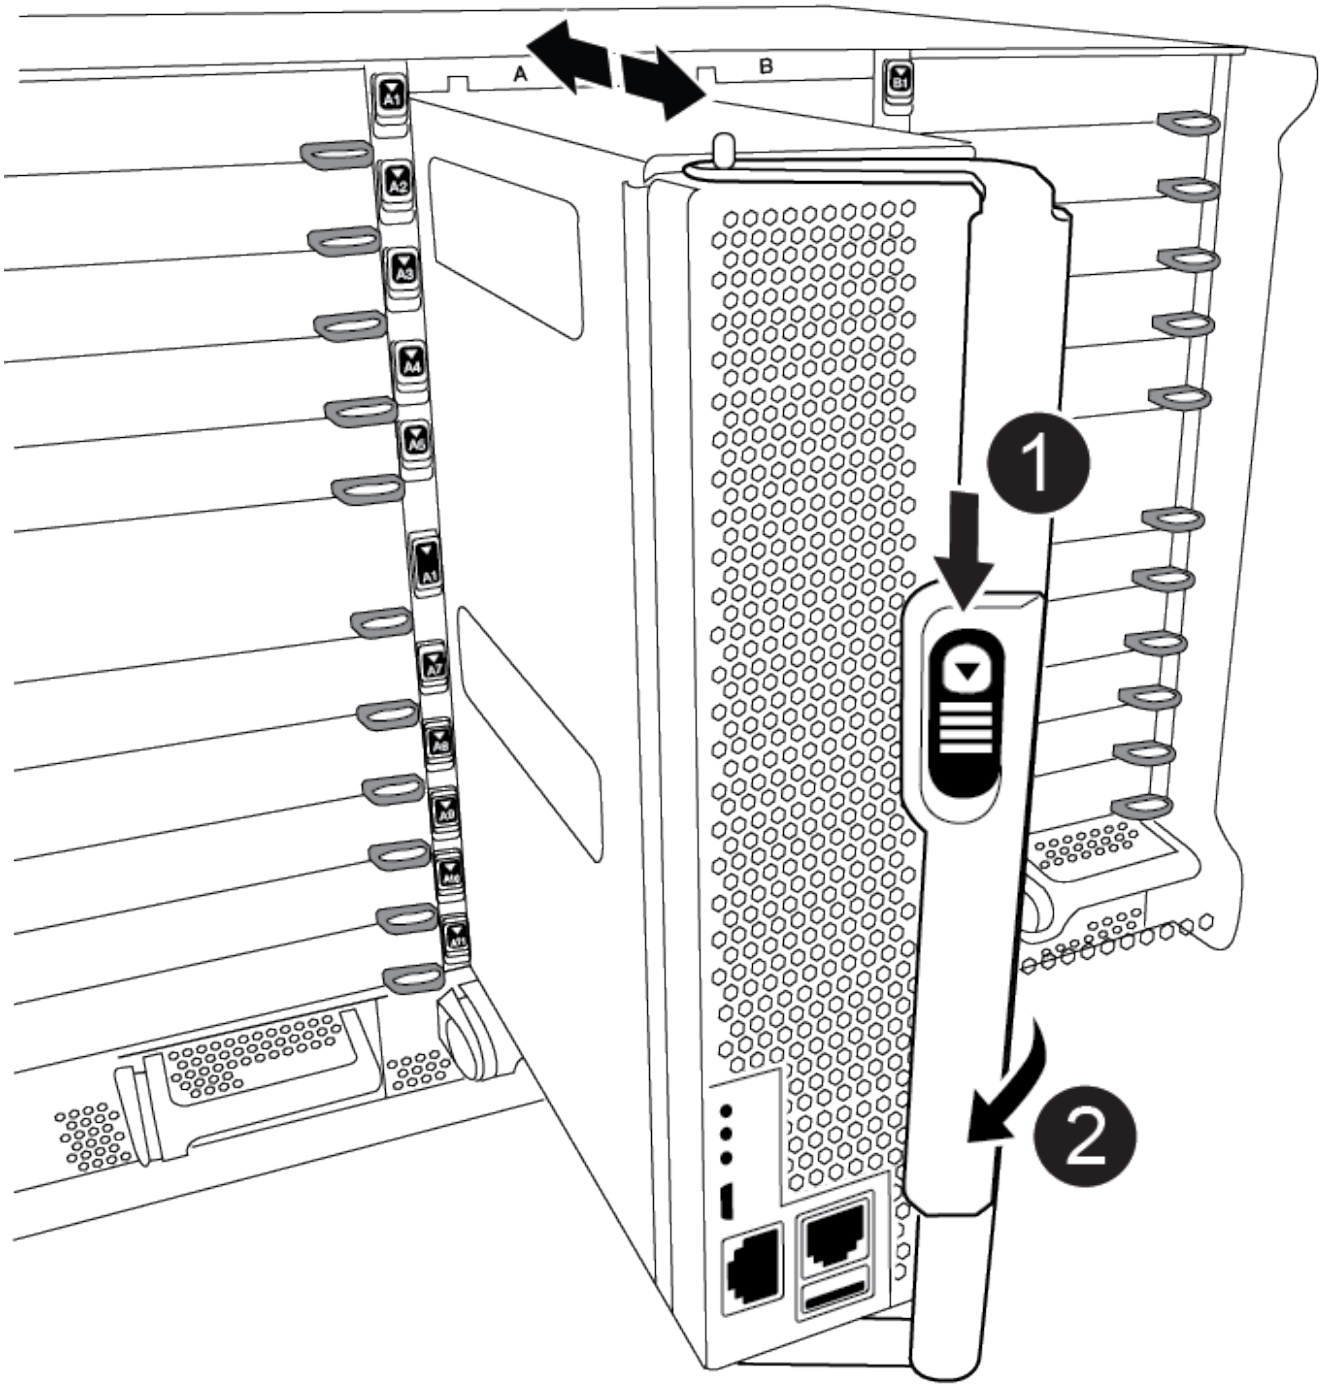

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

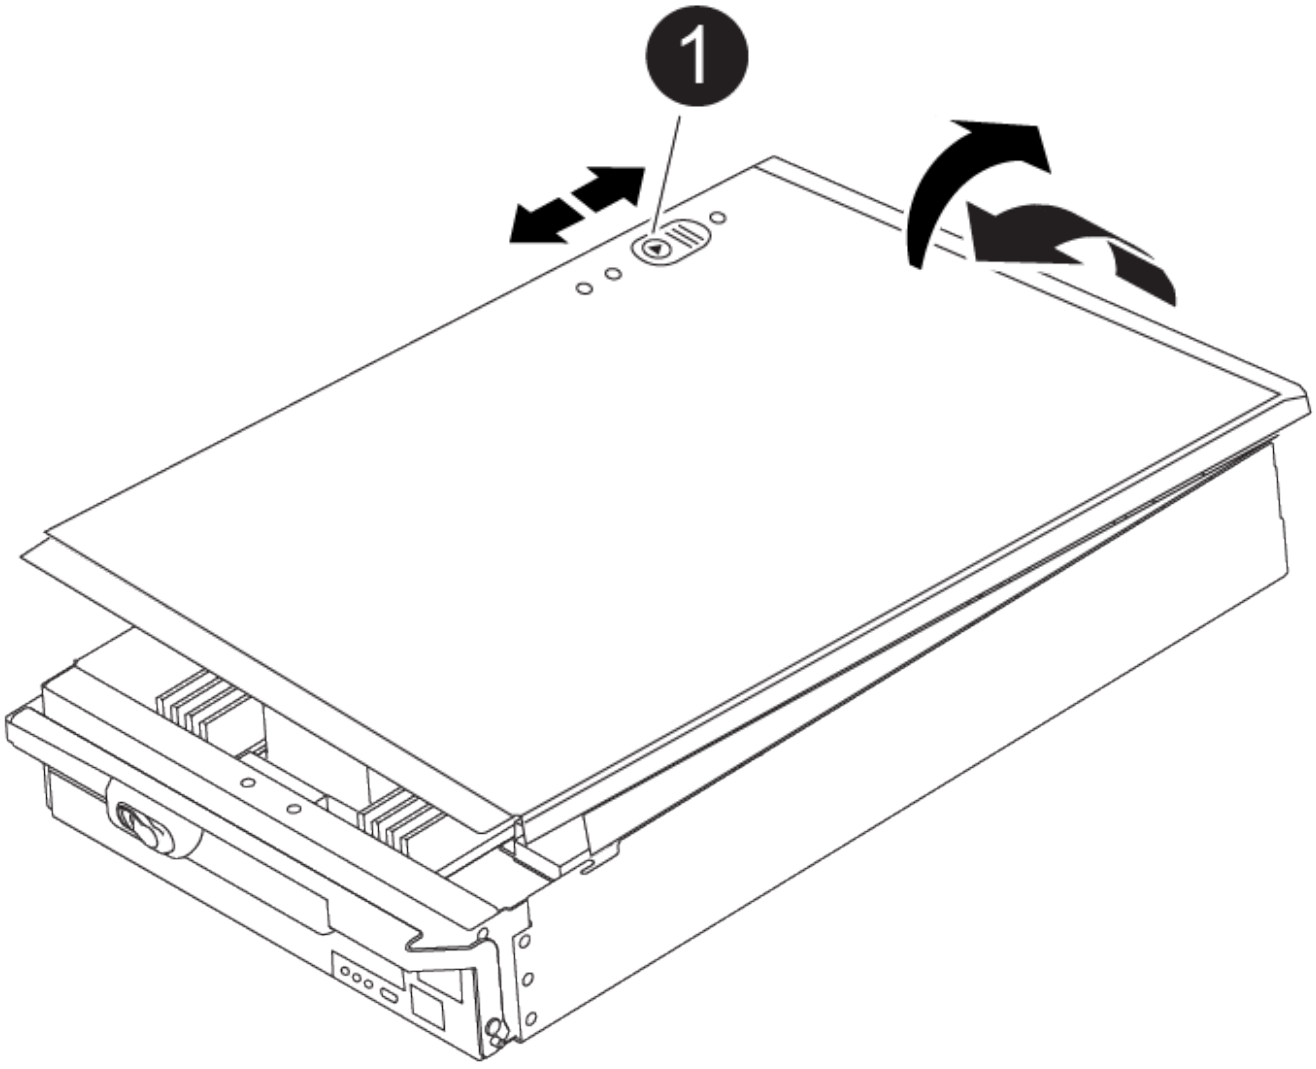

1	Controller module cover locking button

6. Replace the boot media:

a. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

Animation - Replace boot media

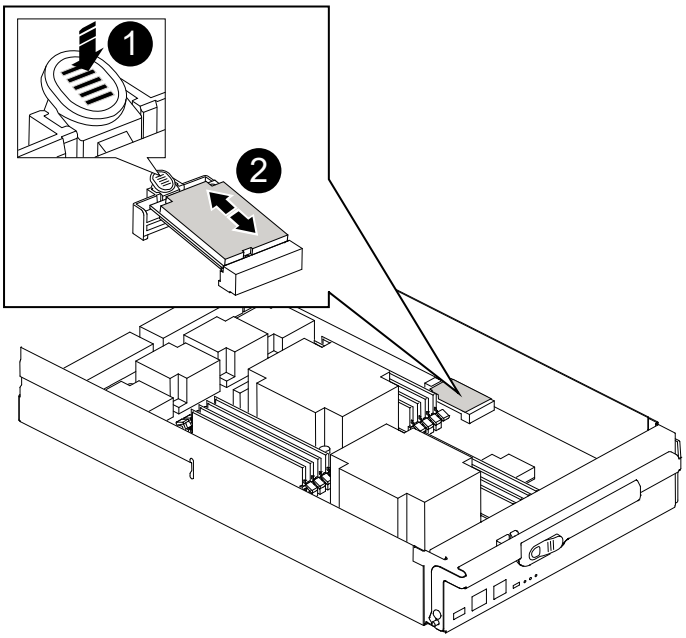

1	Press release tab
2	Boot media

b. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

c. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.

d. Check the boot media to make sure that it is seated squarely and completely in the socket.

If necessary, remove the boot media and reseat it into the socket.

e. Push the boot media down to engage the locking button on the boot media housing.

7. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.
8. Reinstall the controller module:
   a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
   b. Recable the controller module, as needed.

c. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, and then push the cam handle to the closed position.

   The controller begins to boot as soon as it is completely installed into the chassis.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

9. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t iniitator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - AFF A900

After installing the new boot media device in your AFF A900 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - AFF A900

If a component in your AFF A900 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF A900

Get started with replacing the boot media in your AFF A900 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONTAP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - AFF A900

Before replacing the boot media in your AFF A900 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - AFF A900

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - AFF A900

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media and prepare for manual boot recovery - AFF A900

You must unplug the controller module, remove and open the controller module, locate and replace the boot media in the controller, and then transfer the image to the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

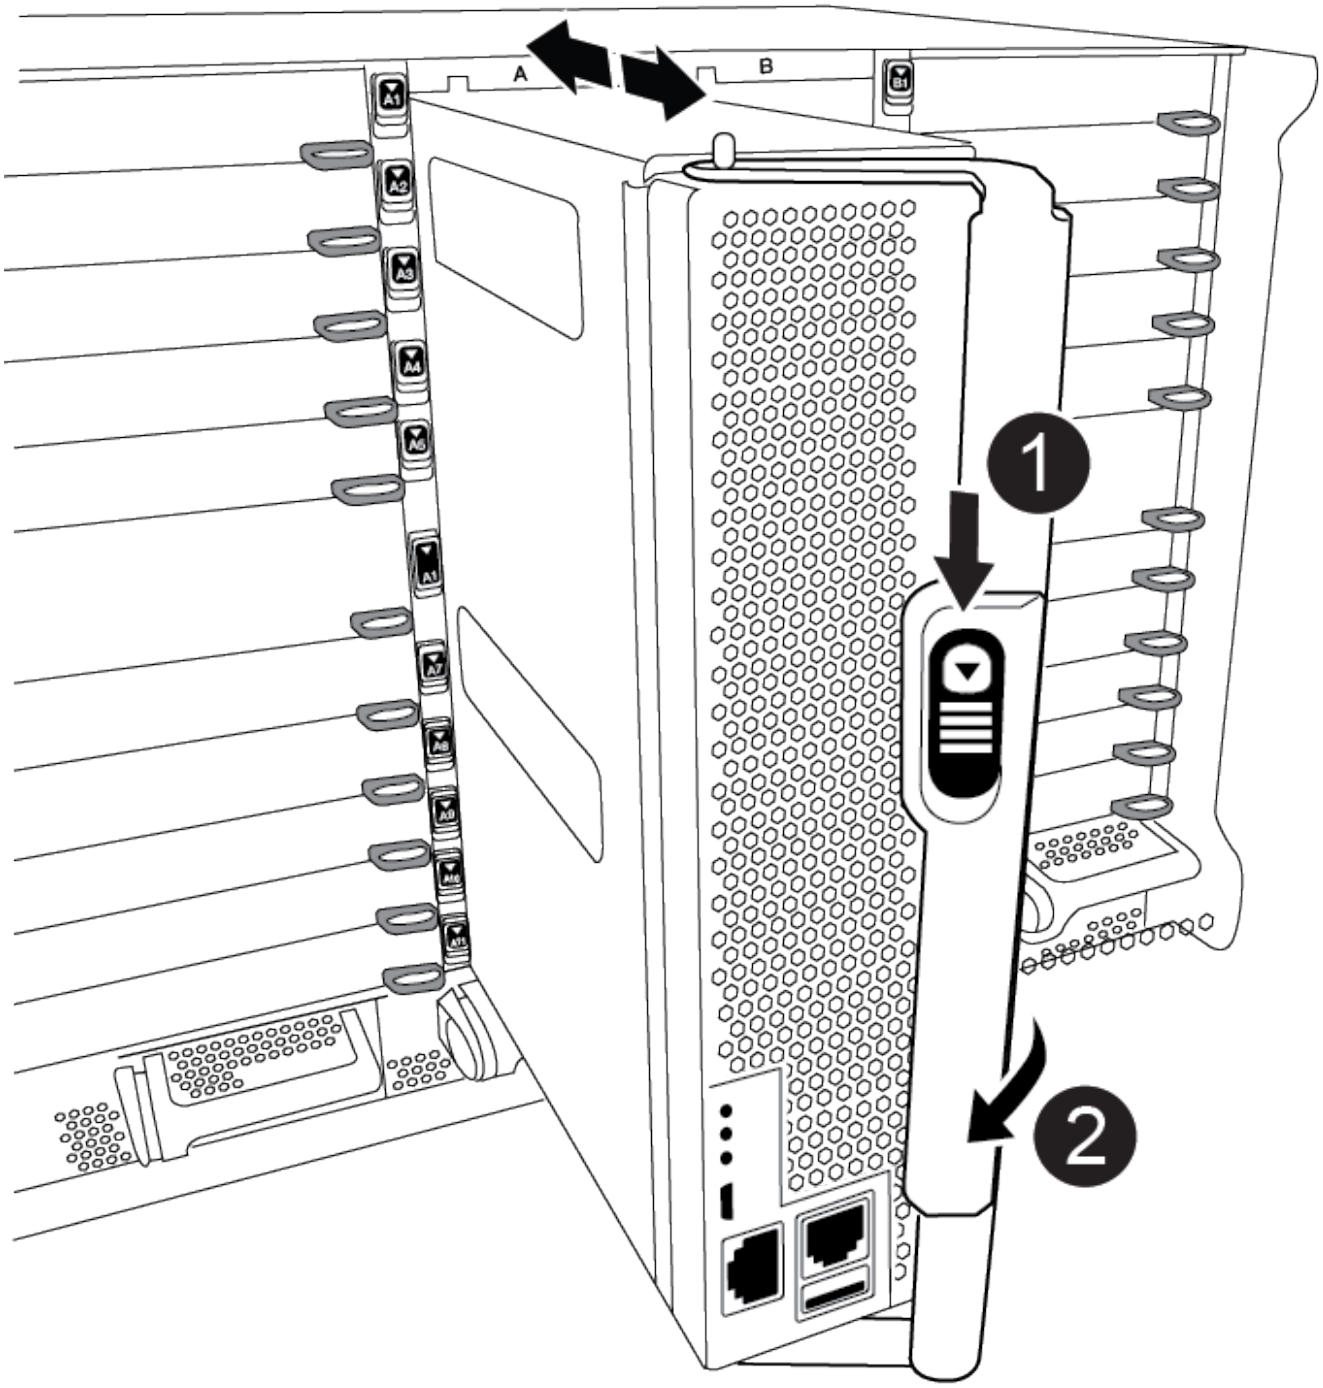

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

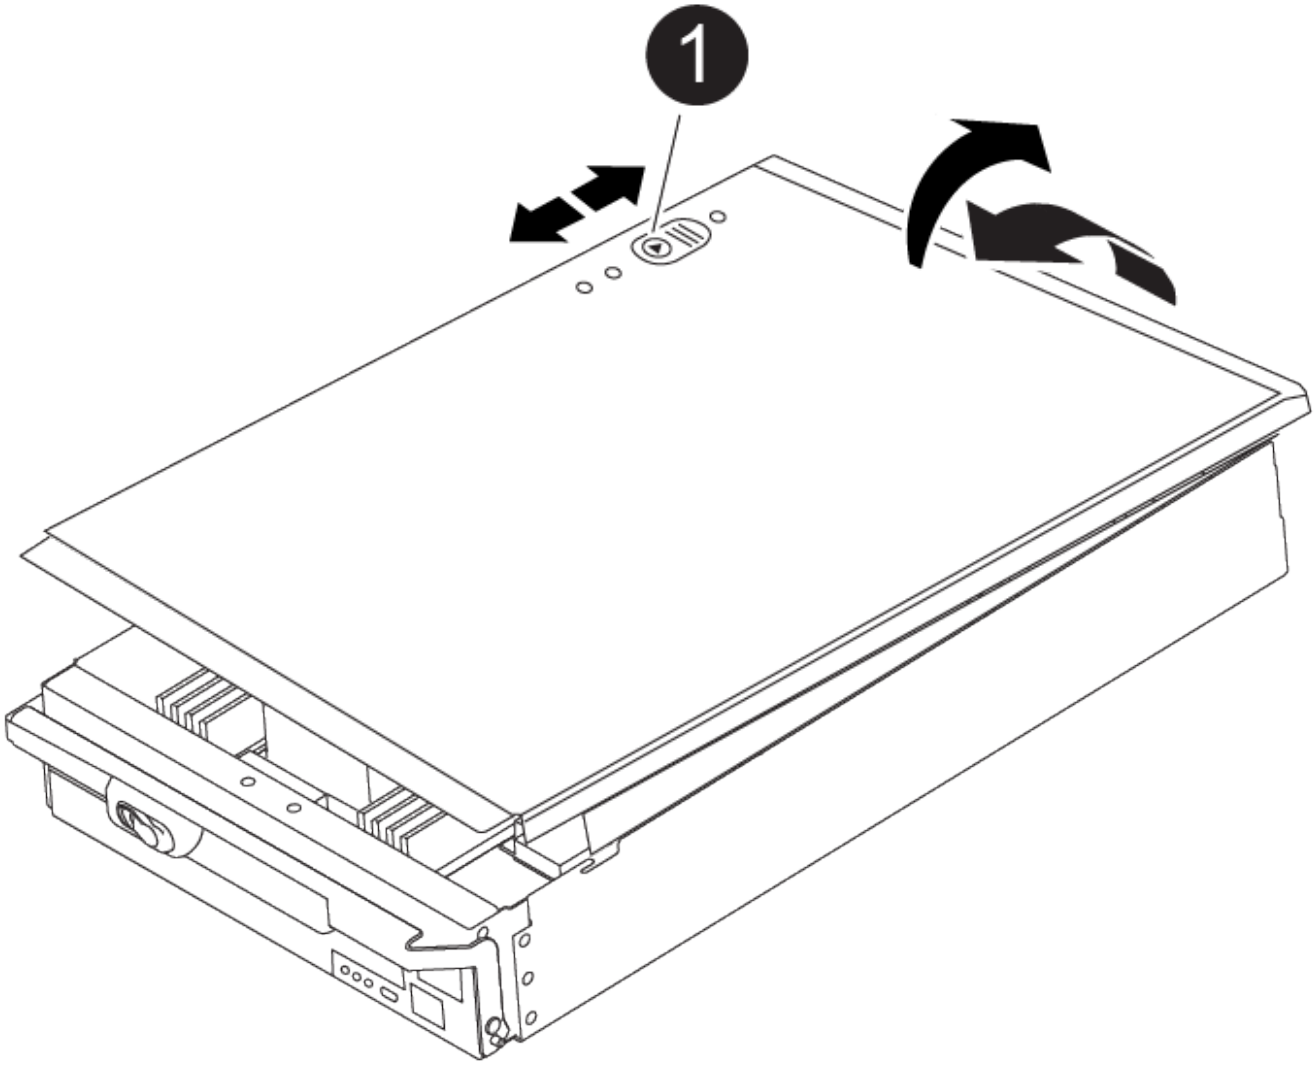

1	Controller module cover locking button

Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

   Animation - Replace boot media

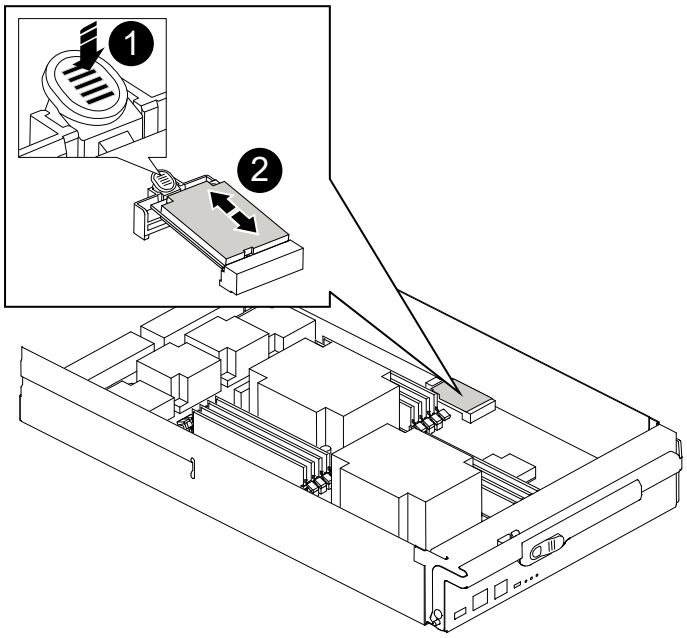

1	Press release tab
2	Boot media

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.
6. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- Download a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site. Use the

`version -v` command to display if your version of ONTAP supports NVE. If the command output displays `<10no- DARE>`, your version of ONTAP does not support NVE.

  - If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. If you have not done so, download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image from the Downloads link on the page, to your work space on your laptop.
   b. Unzip the service image.

      > If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      The USB flash drive should have the appropriate ONTAP image of what the impaired controller is running.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
3. Recable the controller module, as needed.
4. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

5. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, and then push the cam handle to the closed position.

   The controller begins to boot as soon as it is completely installed into the chassis.

6. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

7. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - AFF A900

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore encryption - AFF A900

Restore encryption on the node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*******************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*******************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - AFF A900

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Replace the chassis - AFF A900

To replace the chassis, you must remove the power supplies, fans, controller modules, I/O modules, DCPM modules, and USB LED module from the impaired chassis, remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis in its place, and then install the components into the replacement chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shutdown the controllers - AFF A900

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   ```

```
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - AFF A900

To replace the chassis, you must remove the components from the impaired chassis and install them in the replacement chassis.

Step 1: Remove the power supplies

Removing the power supplies when replacing a chassis involves turning off, disconnecting, and then removing the four power supplies from the rear of the impaired chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Press and hold the terra cotta locking button on the power supply handle, and then pull the power supply out of the chassis.

   > When removing a power supply, always use two hands to support its weight.

   Animation - Remove/install PSU

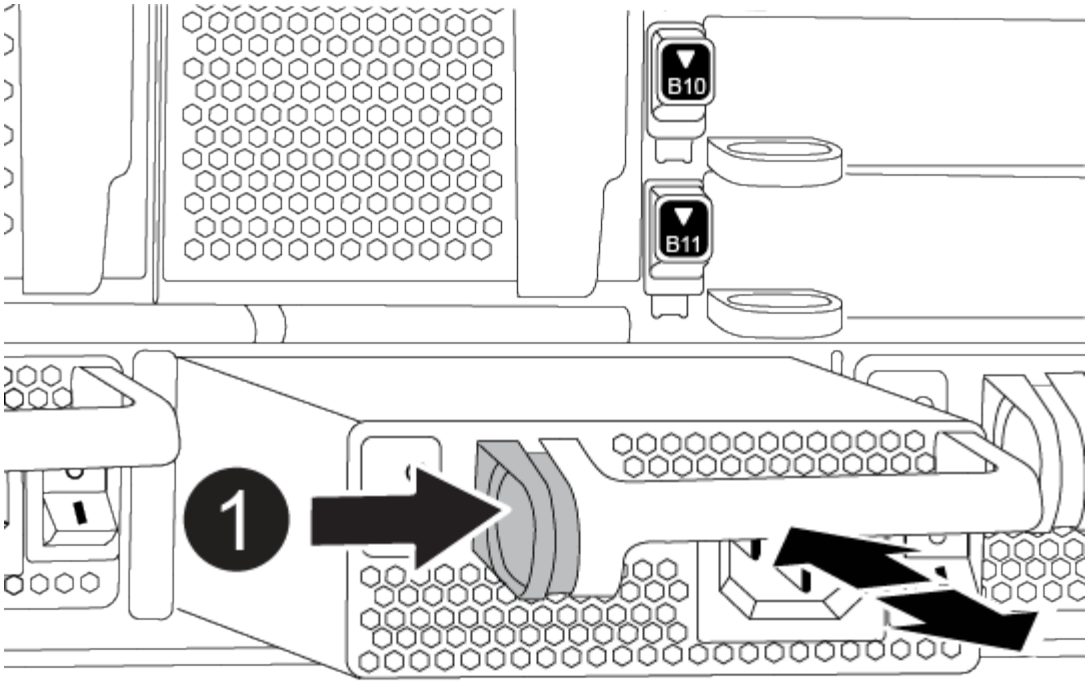

1	Locking button

4. Repeat the preceding steps for any remaining power supplies.

Step 2: Remove the fans

You must remove the six fan modules, located on in the front of the chassis, when replacing the chassis.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press the terra cotta locking button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

> The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

Animation - Remove/install fan

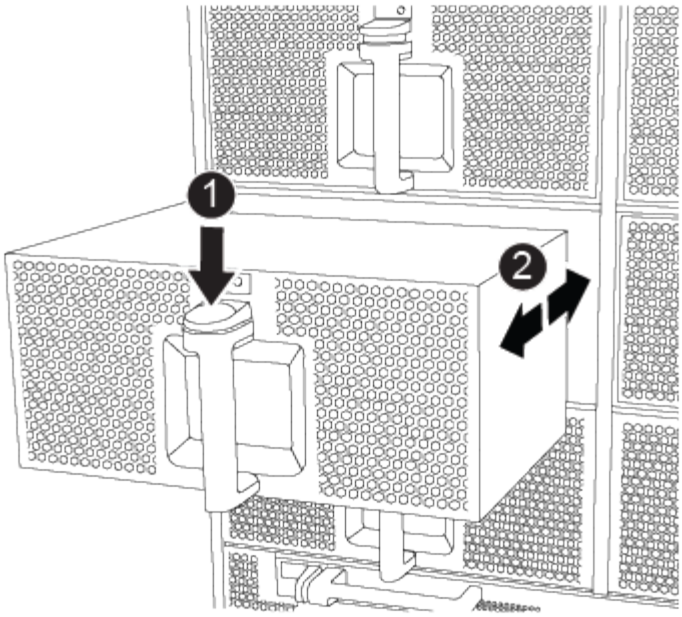

1	Terra cotta locking button
2	Slide fan in/out of chassis

4. Set the fan module aside.
5. Repeat the preceding steps for any remaining fan modules.

Step 3: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the impaired chassis.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta locking button on the cam handle downward until it unlocks.

   Animation - Remove the controller

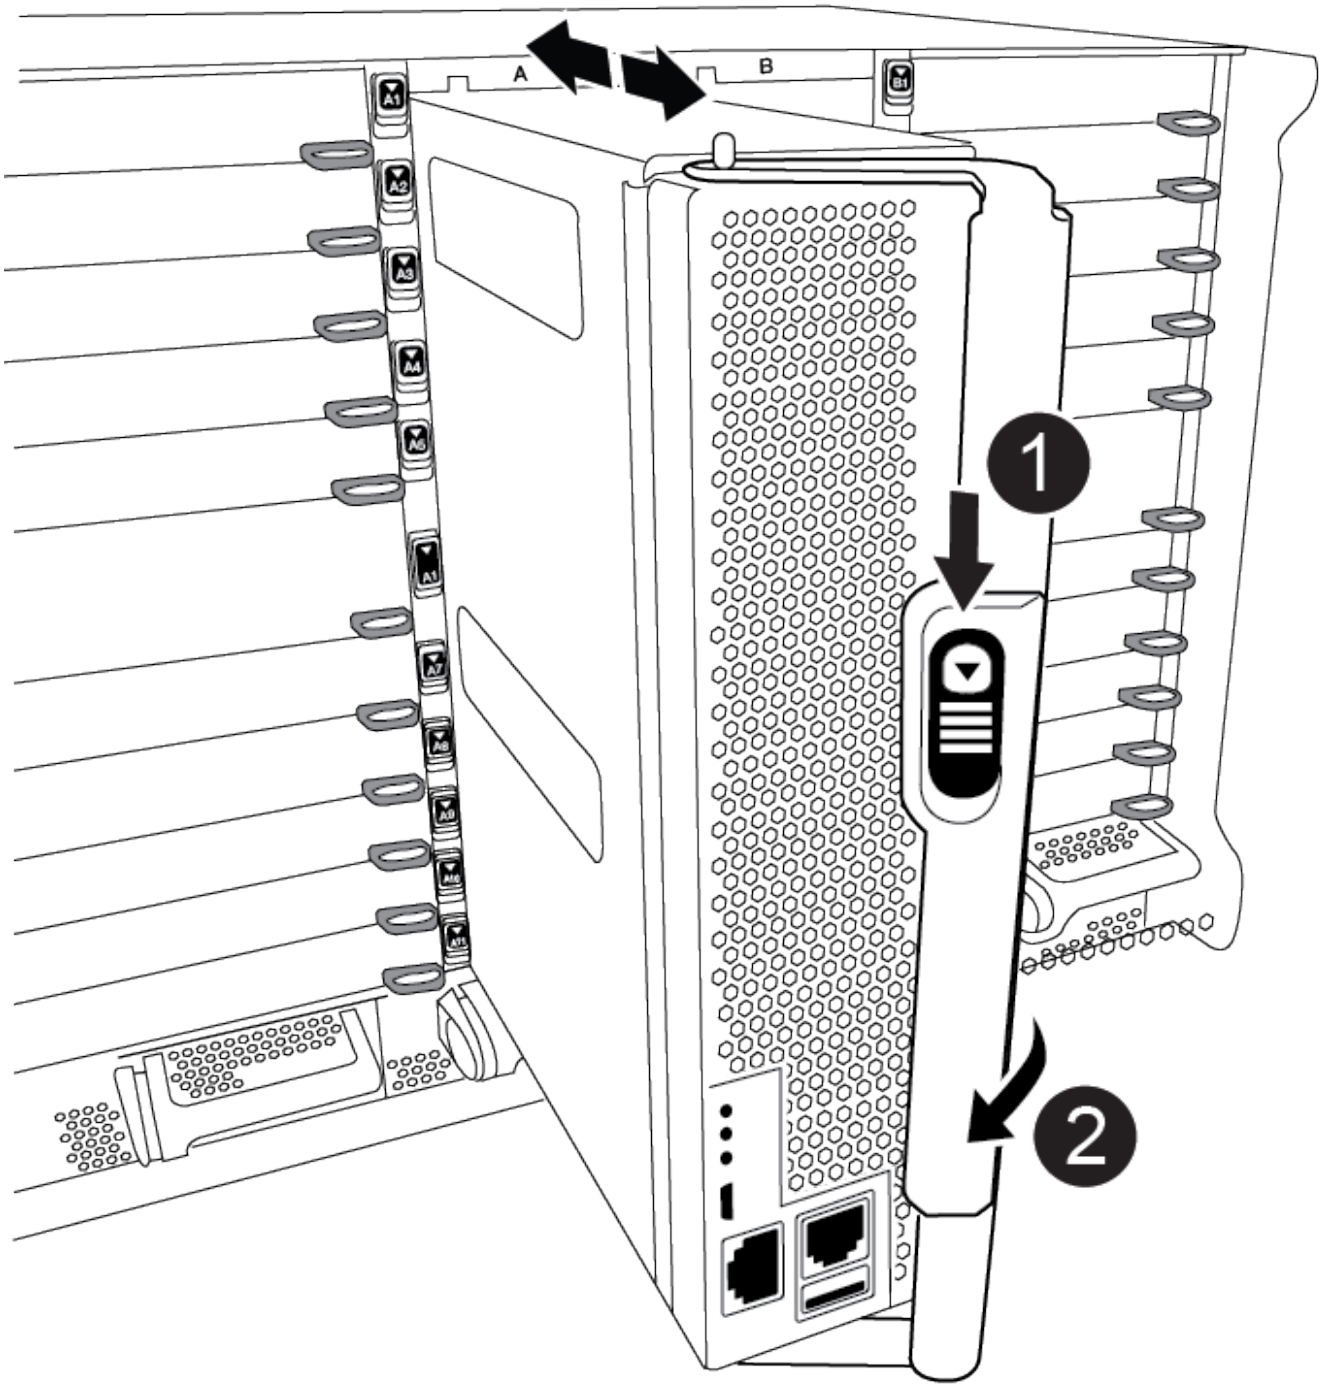

1	Cam handle locking button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Set the controller module aside in a safe place and keep track of which chassis slot it came from, so that it can be installed into the same slot in the replacement chassis..
6. Repeat these steps if you have another controller module in the chassis.

Step 4: Remove the I/O modules

To remove I/O modules from the impaired chassis, including the NVRAM modules, follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

3. Remove the target I/O module from the chassis:

   a. Depress the lettered and numbered cam locking button.

      The cam locking button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

      Animation - Remove/install I/O module

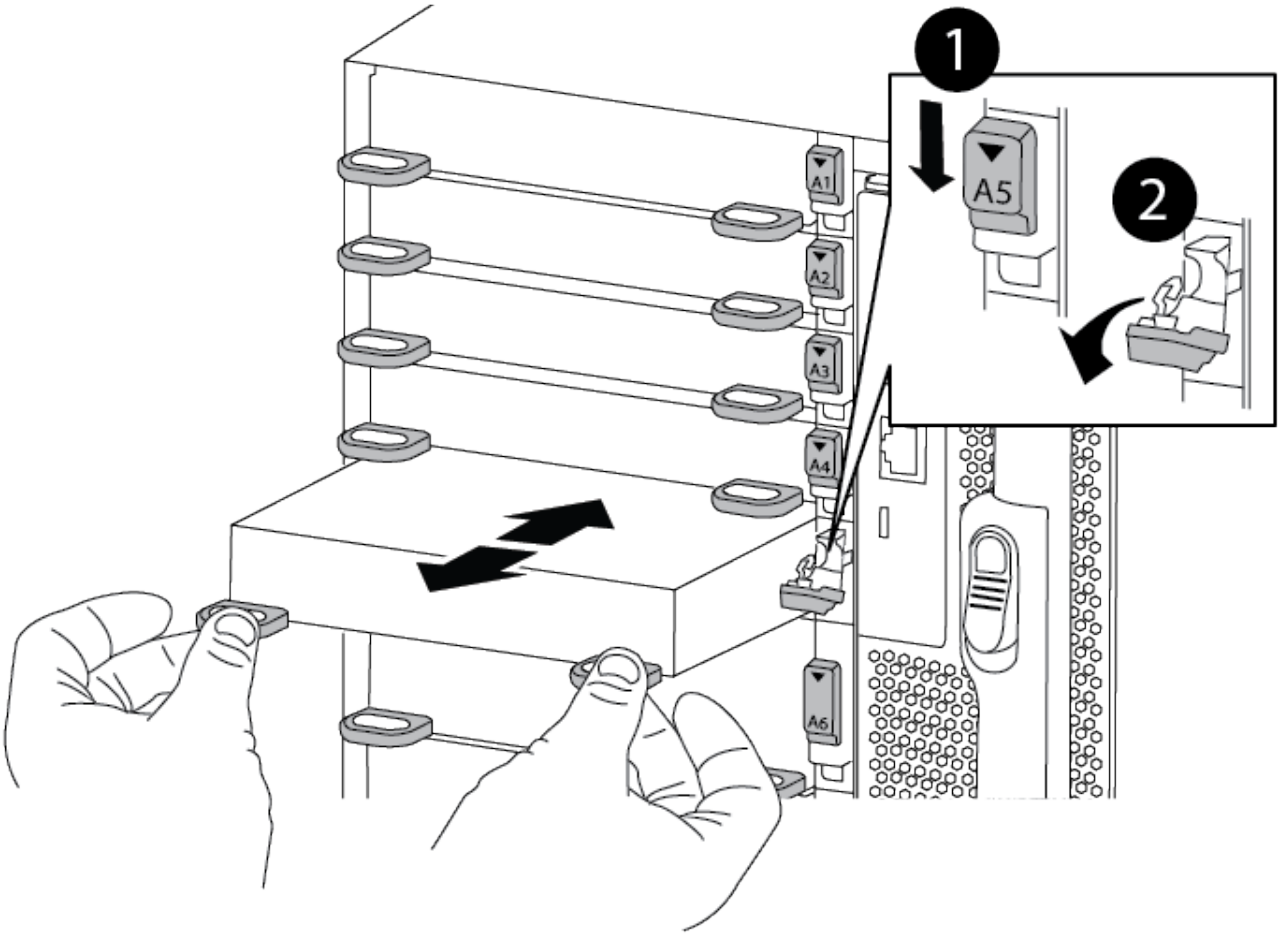

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Repeat the preceding step for the remaining I/O modules in the impaired chassis.

Step 5: Remove the de-stage controller power module

Remove the two de-stage controller power modules from the front of the impaired chassis.

1. If you are not already grounded, properly ground yourself.
2. Press the terra cotta locking button on the module handle, and then slide the DCPM out of the chassis.

   Animation - Remove/install DCPM

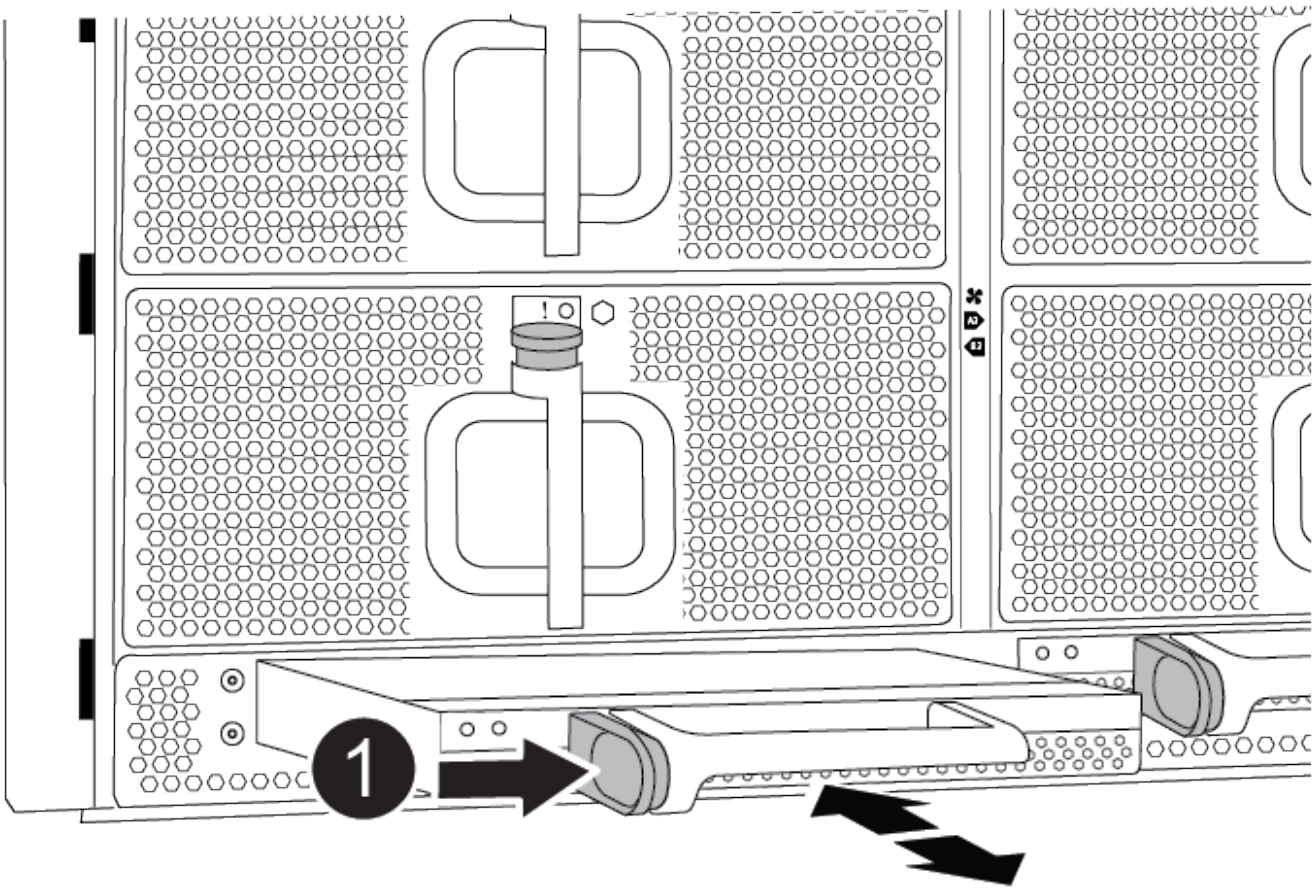

1	DCPM terra cotta locking button

3. Set the DCPM aside in a safe place and repeat this step for the remaining DCPM.

Step 6: Remove the USB LED module

Remove the USB LED modules.

Animation - Remove/install USB

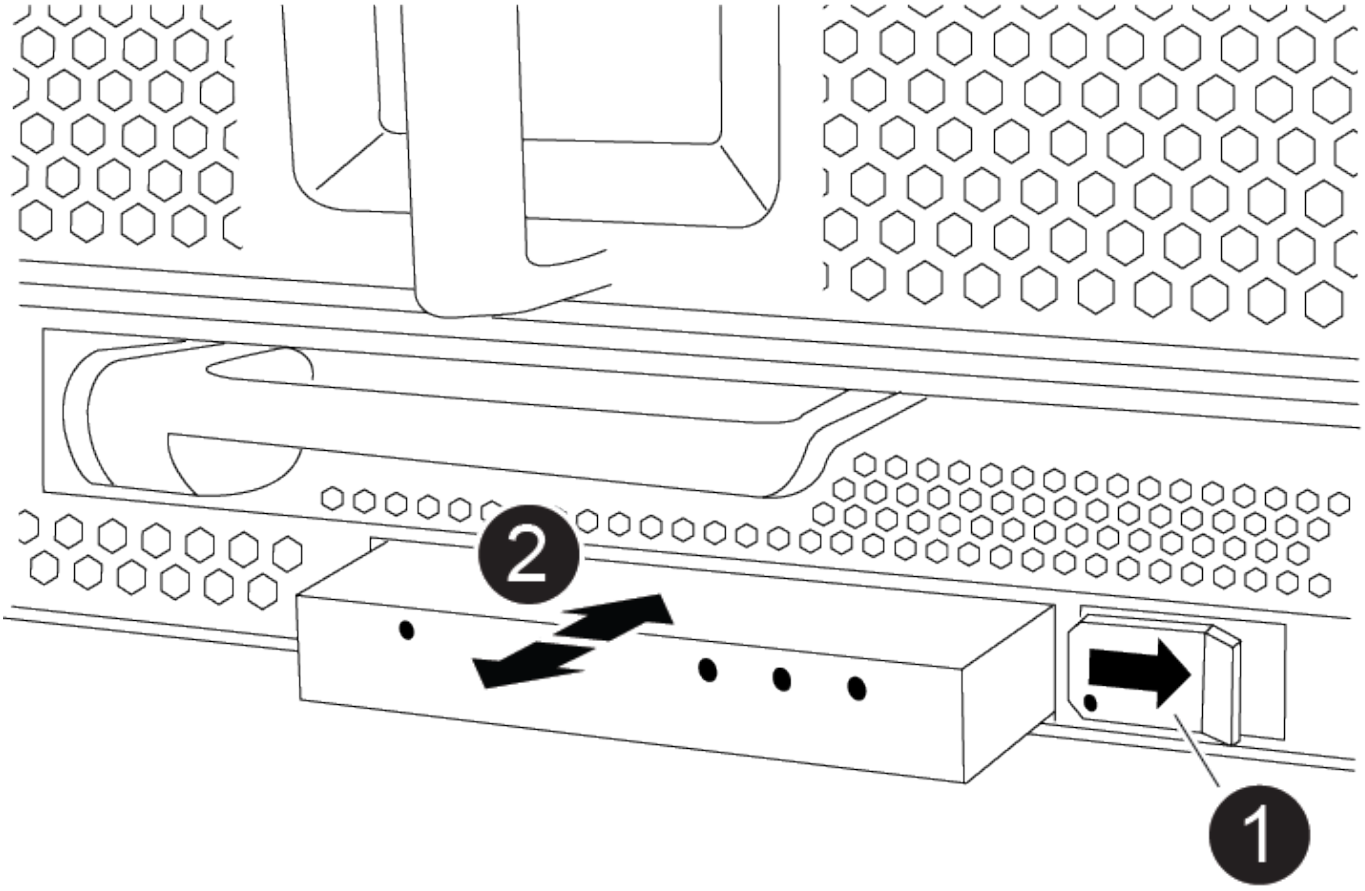

1	Eject the module.
2	Slide out of chassis.

1. Locate the USB LED module on the front of the impaired chassis, directly under the DCPM bays.
2. Press the black locking button on the right side of the module to release the module from the chassis, and then slide it out of the impaired chassis.
3. Set the module aside in a safe place.

Step 7: Remove chassis

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.

   If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

2. With the help of two or three people, slide the impaired chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.

4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.
7. Secure the rear of the chassis to the equipment rack or system cabinet.
8. If you are using the cable management brackets, remove them from the impaired chassis, and then install them on the replacement chassis.

Step 8: Install the de-stage controller power module

When the replacement chassis is installed into the rack or system cabinet, you must reinstall the de-stage controller power modules into it.

1. If you are not already grounded, properly ground yourself.
2. Align the end of the DCPM with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

3. Repeat this step for the remaining DCPM.

Step 9: Install fans into the chassis

To install the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

1. If you are not already grounded, properly ground yourself.
2. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

3. Repeat these steps for the remaining fan modules.
4. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Step 10: Install I/O modules

To install I/O modules, including the NVRAM modules from the impaired chassis, follow the specific sequence of steps.

You must have the chassis installed so that you can install the I/O modules into the corresponding slots in the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. After the replacement chassis is installed in the rack or cabinet, install the I/O modules into their corresponding slots in the replacement chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage, and then push the I/O cam latch all the way up to lock the module in place.

3. Recable the I/O module, as needed.
4. Repeat the preceding step for the remaining I/O modules that you set aside.

> If the impaired chassis has blank I/O panels, move them to the replacement chassis at this time.

Step 11: Install the power supplies

Installing the power supplies when replacing a chassis involves installing the power supplies into the replacement chassis, and connecting to the power source.

1. If you are not already grounded, properly ground yourself.
2. Make sure the power supplies rockers are in the off position.
3. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

> Do not use excessive force when sliding the power supply into the system. You can damage the connector.

4. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

> Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

5. Repeat the preceding steps for any remaining power supplies.

Step 12: Install the USB LED modules

Install the USB LED modules in the replacement chassis.

1. Locate the USB LED module slot on the front of the replacement chassis, directly under the DCPM bays.
2. Align the edges of the module with the USB LED bay, and gently push the module all the way into the chassis until it clicks into place.

Step 13: Install the controller

After you install the controller module and any other components into the replacement chassis, boot it.

1. If you are not already grounded, properly ground yourself.
2. Connect the power supplies to different power sources, and then turn them on.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

> Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the console to the controller module, and then reconnect the management port.
5. With the cam handle in the open position, slide the controller module into the chassis and firmly push the

controller module in until it meets the midplane and is fully seated, and then close the cam handle until it clicks into the locked position.

> Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.

The controller module begins to boot as soon as it is fully seated in the chassis.

6. Repeat the preceding steps to install the second controller into the replacement chassis.
7. Boot each controller.

Restore and verify the configuration - AFF A900

To complete the chassis replacement, you must complete specific tasks.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`ha-state`*

      The value for HA-state can be one of the following:

      - ha
      - non-ha
3. Confirm that the setting has changed: `ha-config show`
4. If you have not already done so, recable the rest of your system.

Step 2: Bring up the system

1. If you have not done so, plug the power cables back into the PSUs.
2. Turn on the PSUs by toggling the rocker switched to **ON**, and wait for the controllers to power up completely.
3. Check the front and the back of the chassis and controllers for any fault lights after power up.
4. Connect to the SP or BMC IP address of the nodes via SSH. This will be the same address used to shut down the nodes.
5. Perform additional health checks as described in How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP
6. Turn AutoSupport back on (end the maintenance window message):
   `system node autosupport invoke -node * -type all -message MAINT=end`

As a best practice, you should do the following:

- Resolve any Active IQ Wellness Alerts and Risks (Active IQ will take time to process post-power up AutoSupports - expect a delay in results)
- Run Active IQ Config Advisor
- Check system health using How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Replace the controller module - AFF A900

To replace the impaired controller module, you must shut down the impaired controller, move the internal components to the replacement controller module, install the replacement controller module, and reboot the replacement controller.

Before you begin

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system has a V_StorageAttach license, you must refer to the additional required steps before performing this procedure.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the replacement controller so that the replacement controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The impaired controller is the controller that is being replaced.

    - The replacement controller is the new controller that is replacing the impaired controller.
    - The healthy controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - AFF A900

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the controller module hardware - AFF A900

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

The following animation shows the whole process of moving components from the impaired to the replacement controller.

Animation - Move components to replacement controller

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

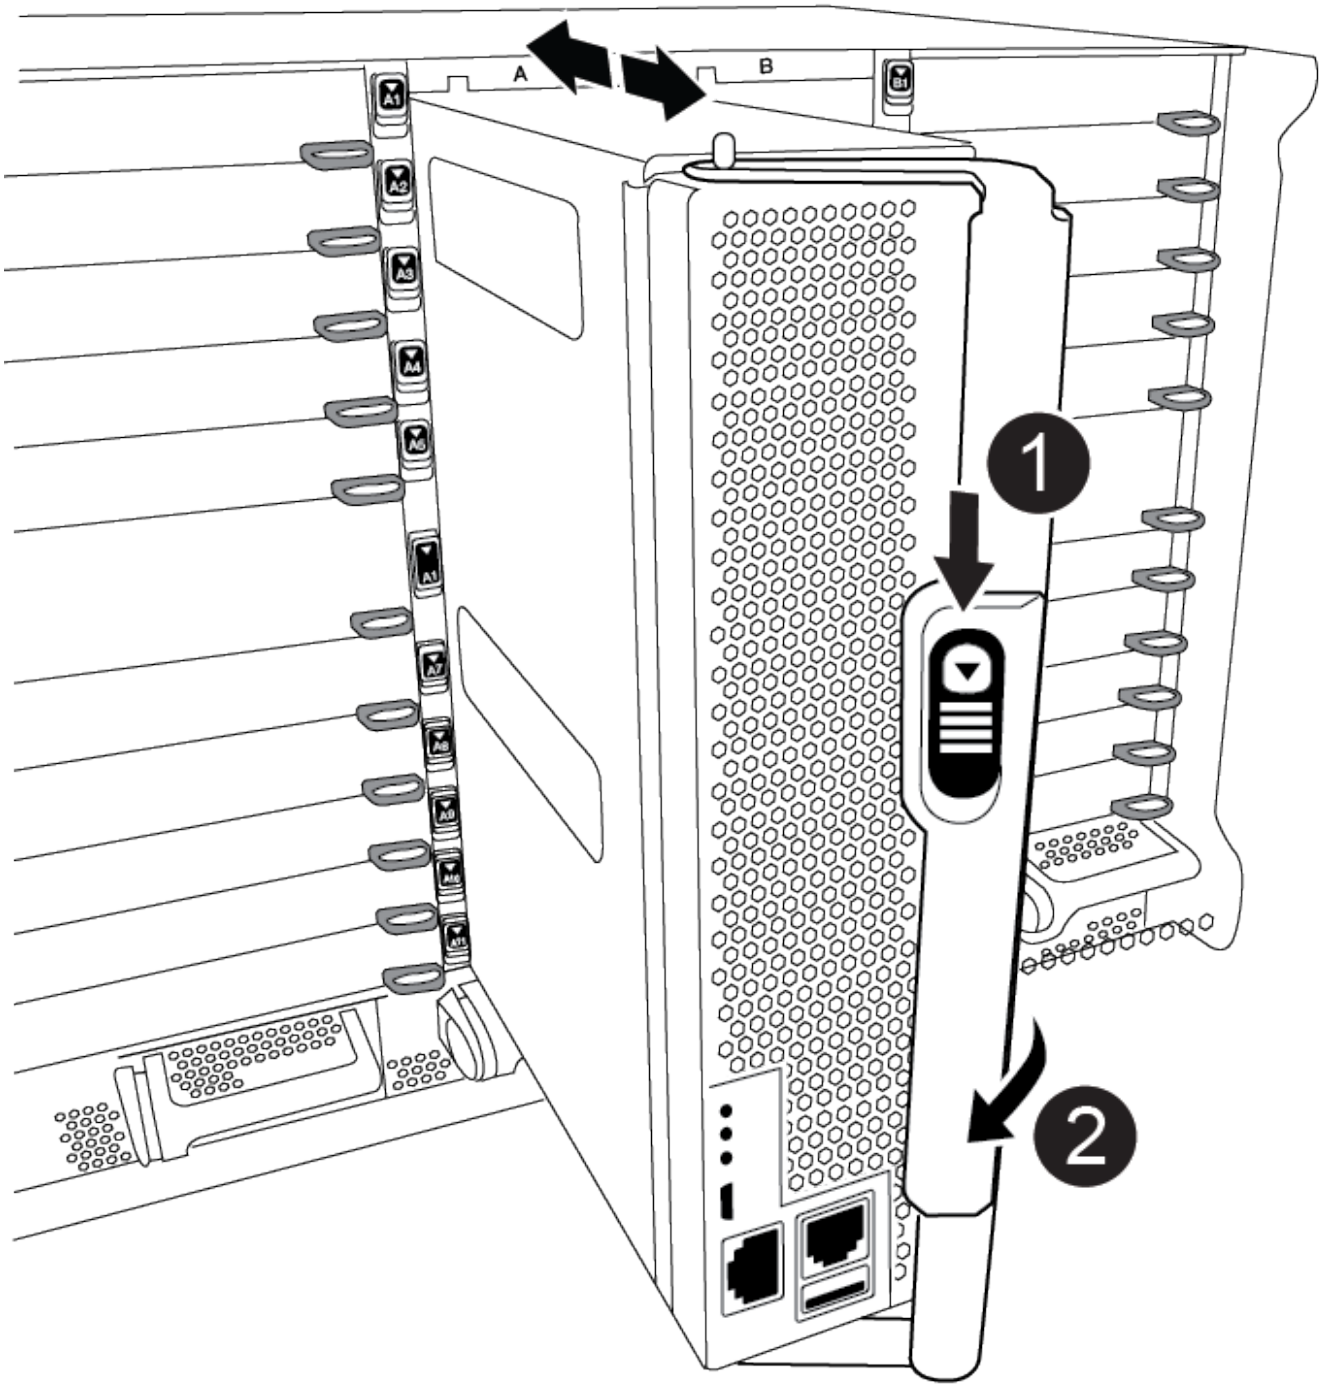

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

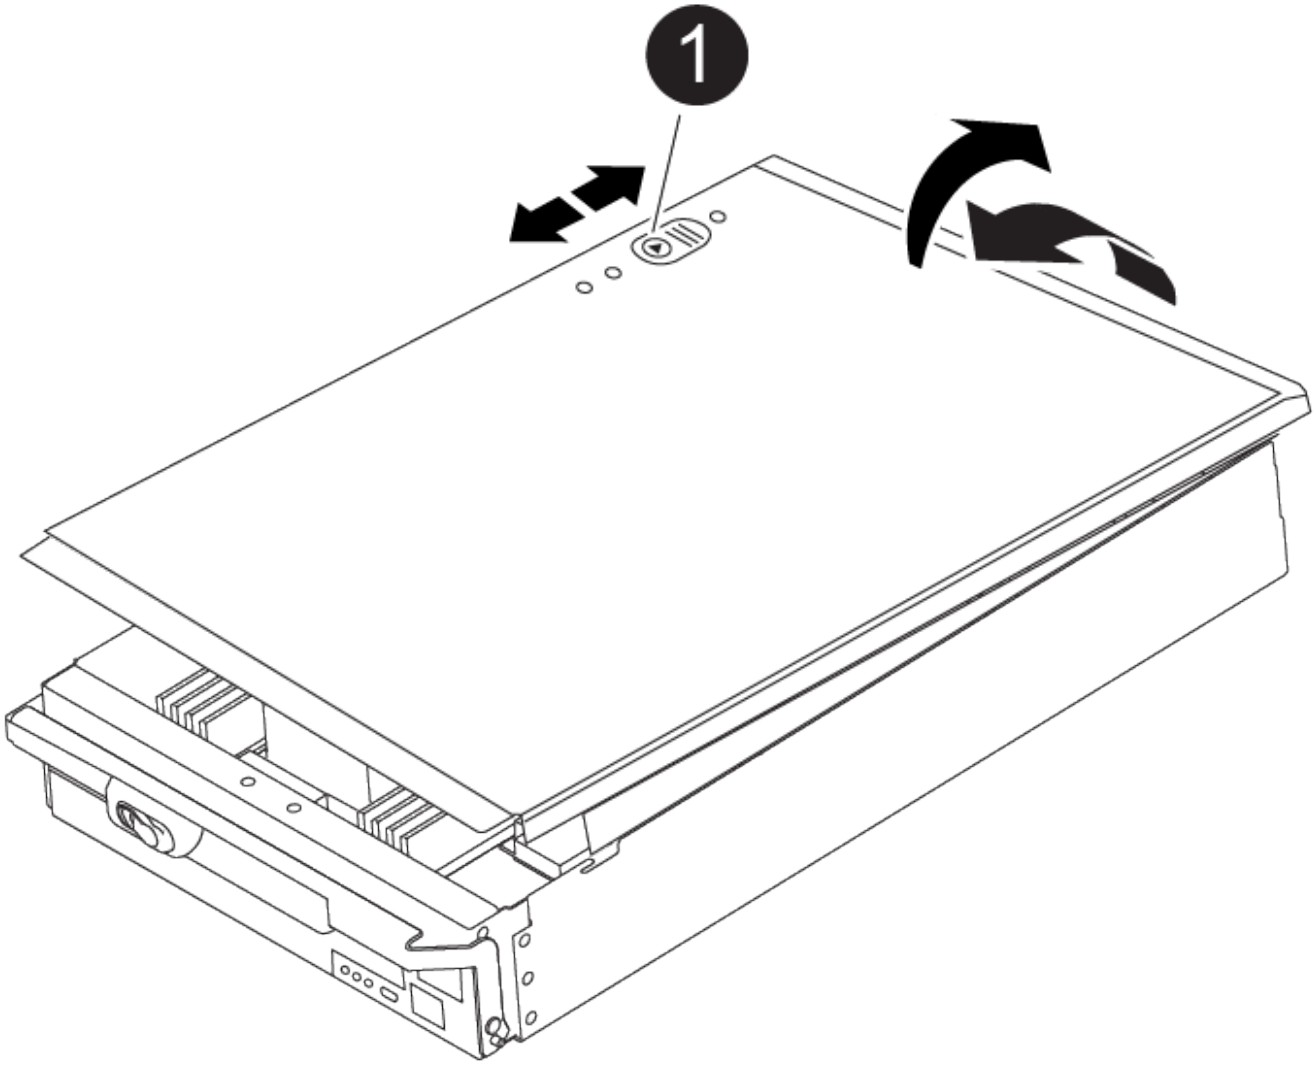

1	Controller module cover locking button

Step 2: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller and insert it in the new controller.

1. Locate the boot media using the following illustration or the FRU map on the controller module:

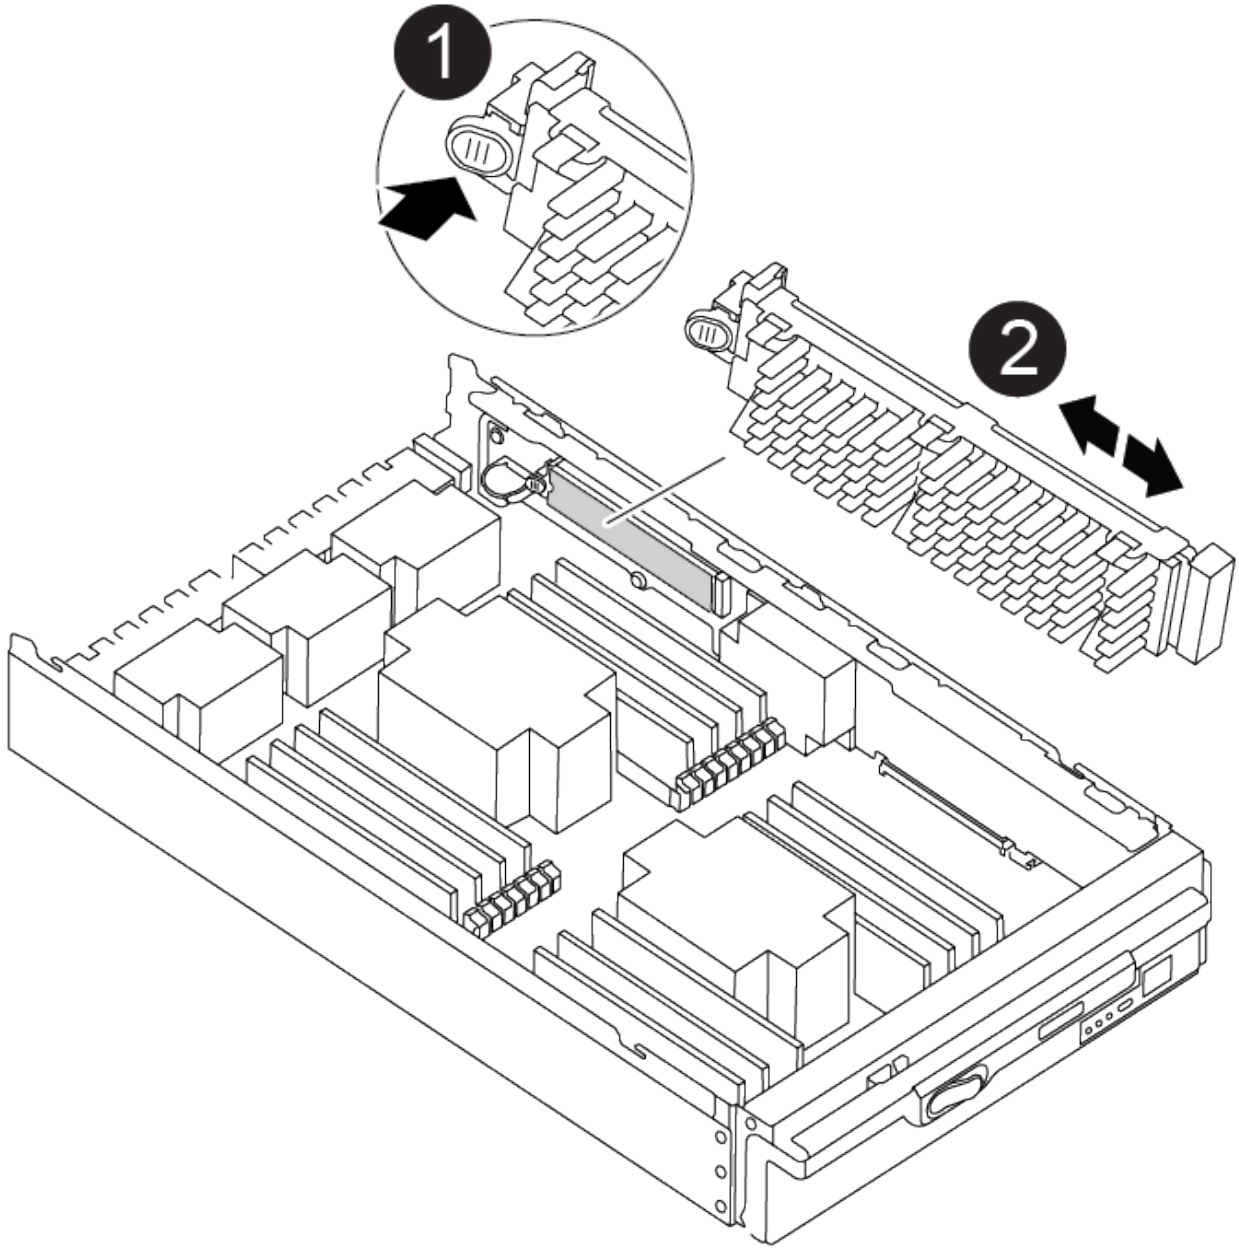

1	Press release tab
2	Boot media

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the system DIMMs

To move the DIMMs, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

> The VER2 controller has fewer DIMM sockets. There is no reduction in the number of DIMMs supported or change in the DIMM socket numbering. When moving the DIMMs to the new controller module, install the DIMMs into the same socket number/location as the impaired controller module. See the FRU map diagram on the VER2 controller module for DIMM socket locations.

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

> Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

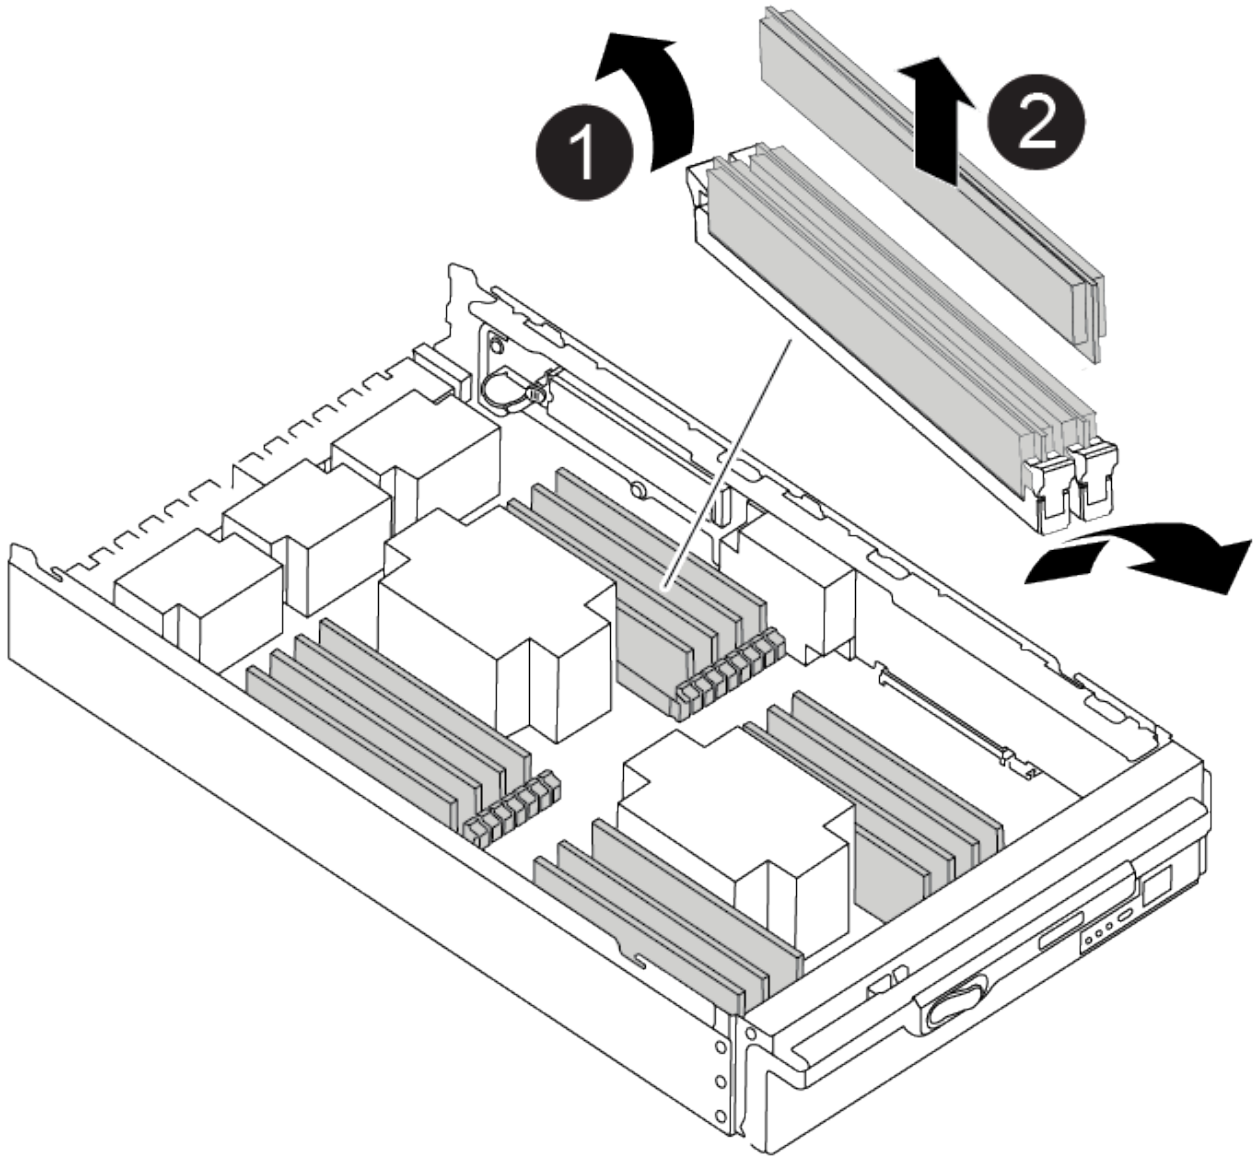

1	DIMM ejector tabs
2	DIMM

5. Locate the slot where you are installing the DIMM.
6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
9. Repeat these steps for the remaining DIMMs.

Step 4: Install the controller

After you install the components into the replacement controller module, you must install the replacement controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Animation - Install controller

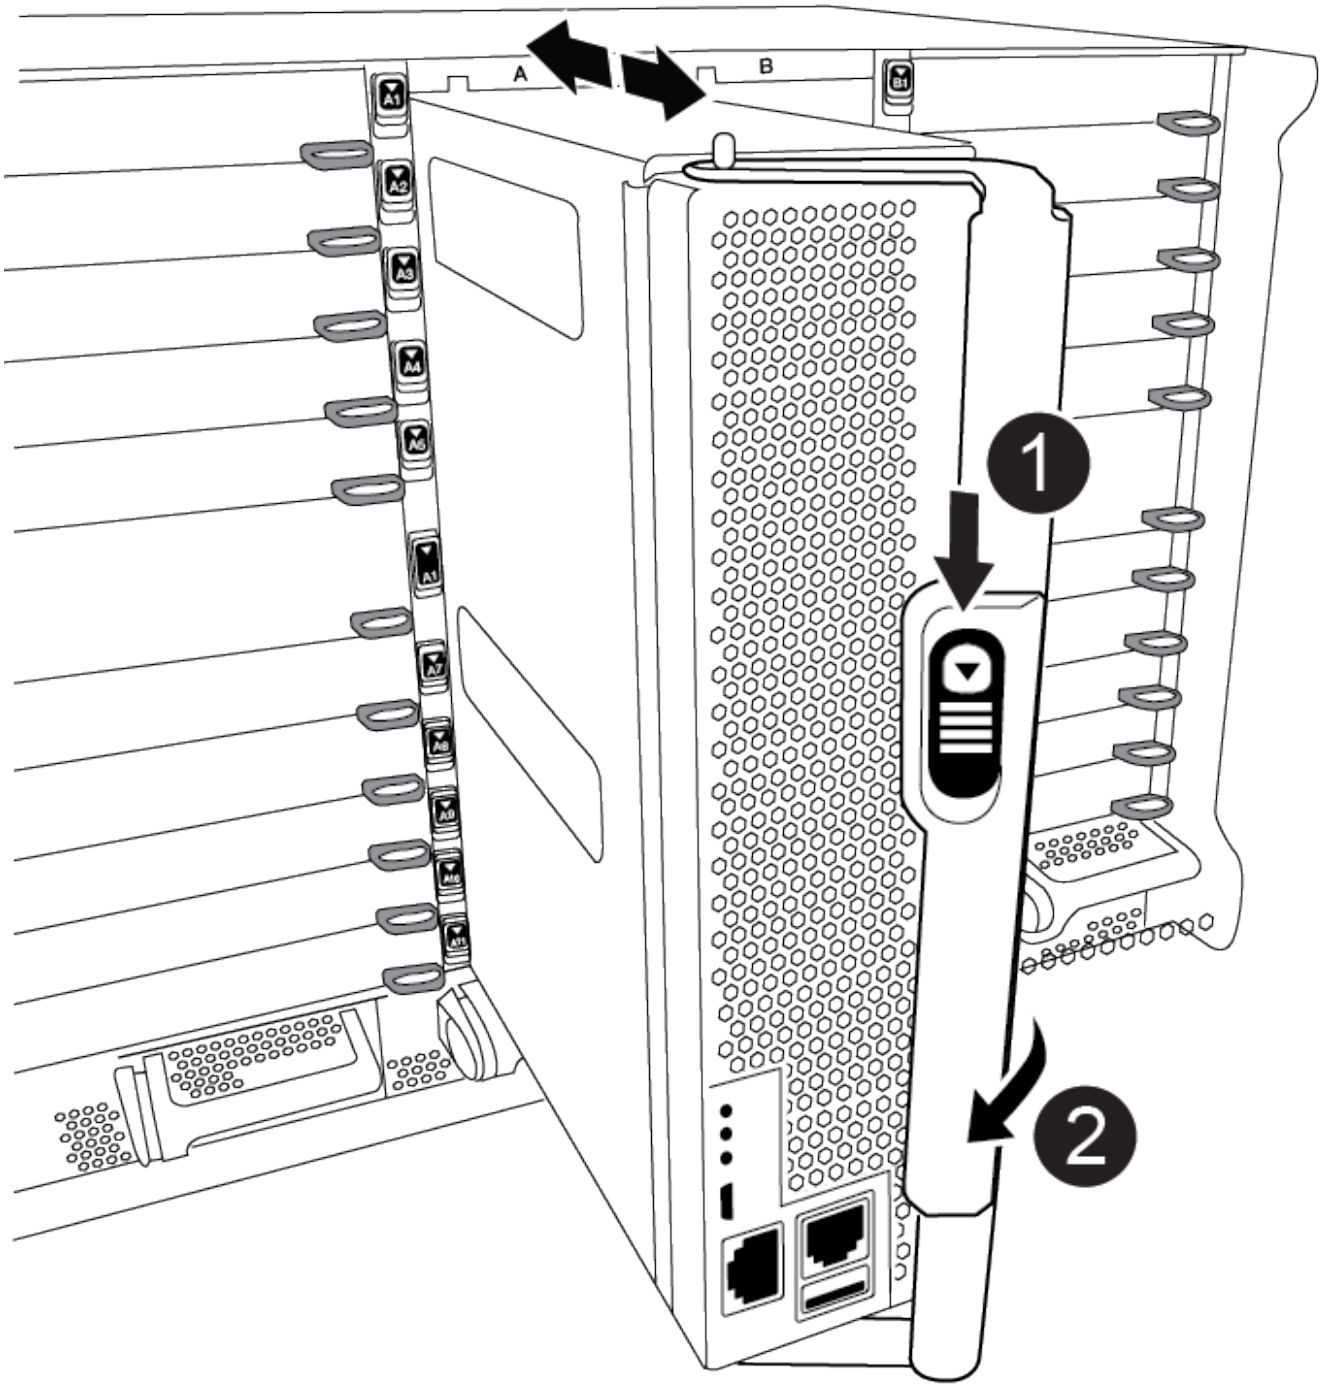

1	Cam handle release button
2	Cam handle

Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in

the following sections.

> You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   c. Rotate the controller module cam handle to the locked position.
   d. Interrupt the boot process by pressing `Ctrl-C` when you see Press Ctrl-C for Boot Menu.
   e. Select the option to boot to LOADER.

Restore and verify the system configuration - AFF A900

After completing the hardware replacement, you verify the low-level system configuration of the replacement controller, and reconfigure system settings as necessary.

Step 1: Set and verify the system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*

6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the controller's HA state

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the replacement controller module, verify that all components display the same `HA` state: `ha-config show`

If your system is in…	The HA state for all components should be…
An HA pair	ha
A MetroCluster FC configuration with four or more nodes	mcc
A MetroCluster IP configuration	mccip

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller` *`ha-state`*
3. If the displayed system state of the chassis does not match your system configuration, set the `HA` state for the chassis: `ha-config modify chassis` *`ha-state`*

Recable the system - AFF A900

Continue the replacement procedure by recabling the storage and network conigurations.

Step 1: Recable the system

You must recable the controller module's storage and network connections.

Steps

1. Recable the system.
2. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

> The system ID and disk assignment information reside in the NVRAM module, which is in a module separate from the controller module and not impacted by the controller module replacement.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`
   c. Wait for the savecore command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *`local-node-name`* `partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

- Restore onboard key management encryption keys
- Restore external key management encryption keys

6. Give back the controller:

   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The 'metrocluster node show -fields node-systemid' command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR

home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

For more information, see Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration topic.

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF A900

To complete the replacement procedure and restore your system to full operation, restore the NetApp Storage Encryption configuration (if necessary) and install licenses for the new controller.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the

feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The licenses keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.
- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

   The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A900

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

Before you begin

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

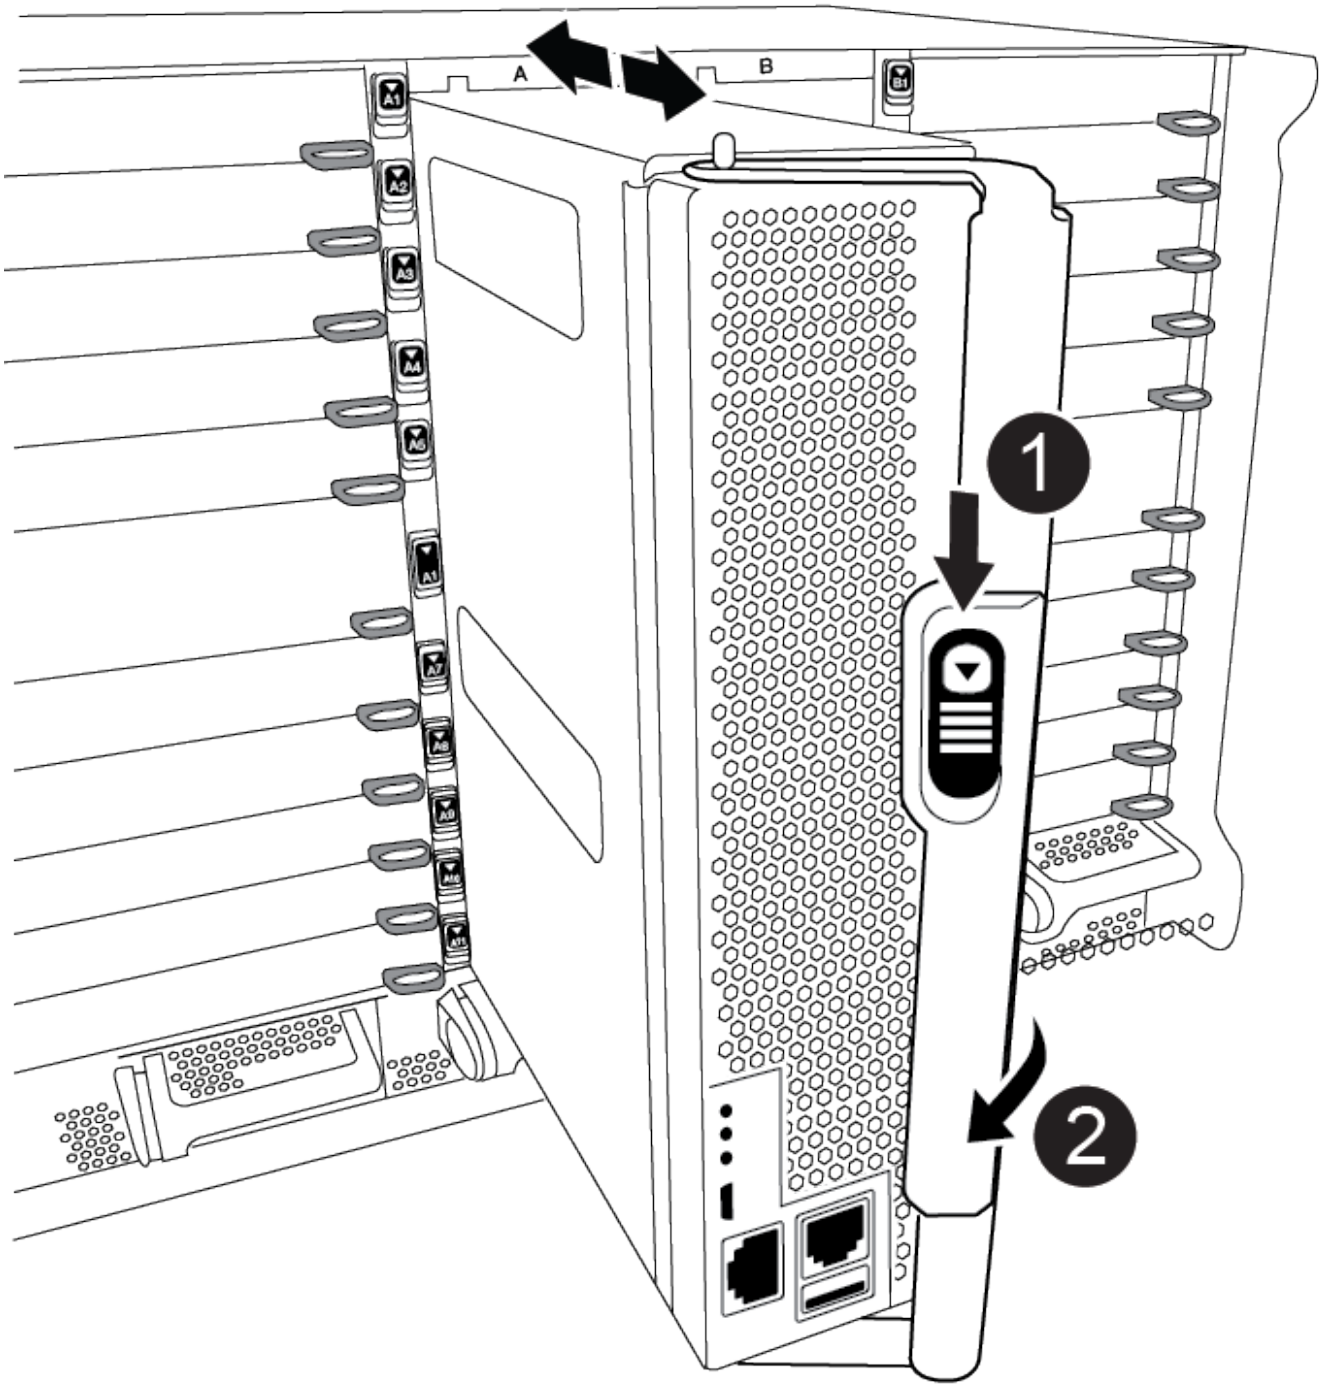

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

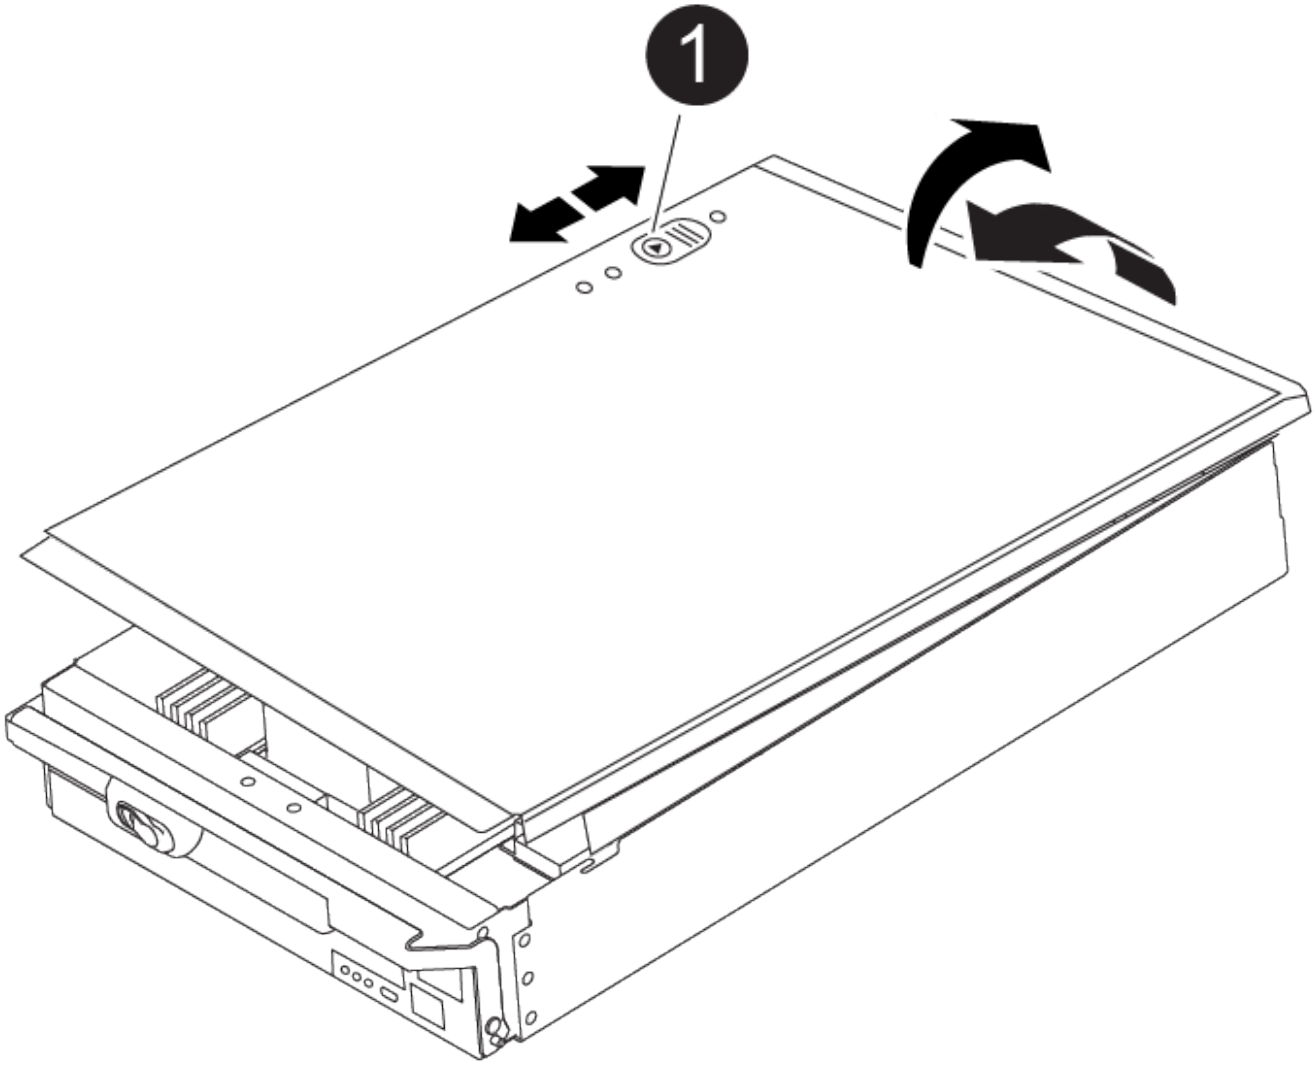

1	Controller module cover locking button

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

> The VER2 controller has fewer DIMM sockets. There is no reduction in the number of DIMMs supported or change in the DIMM socket numbering. When moving the DIMMs to the new controller module, install the DIMMs into the same socket number/location as the impaired controller module. See the FRU map diagram on the VER2 controller module for DIMM socket locations.

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.

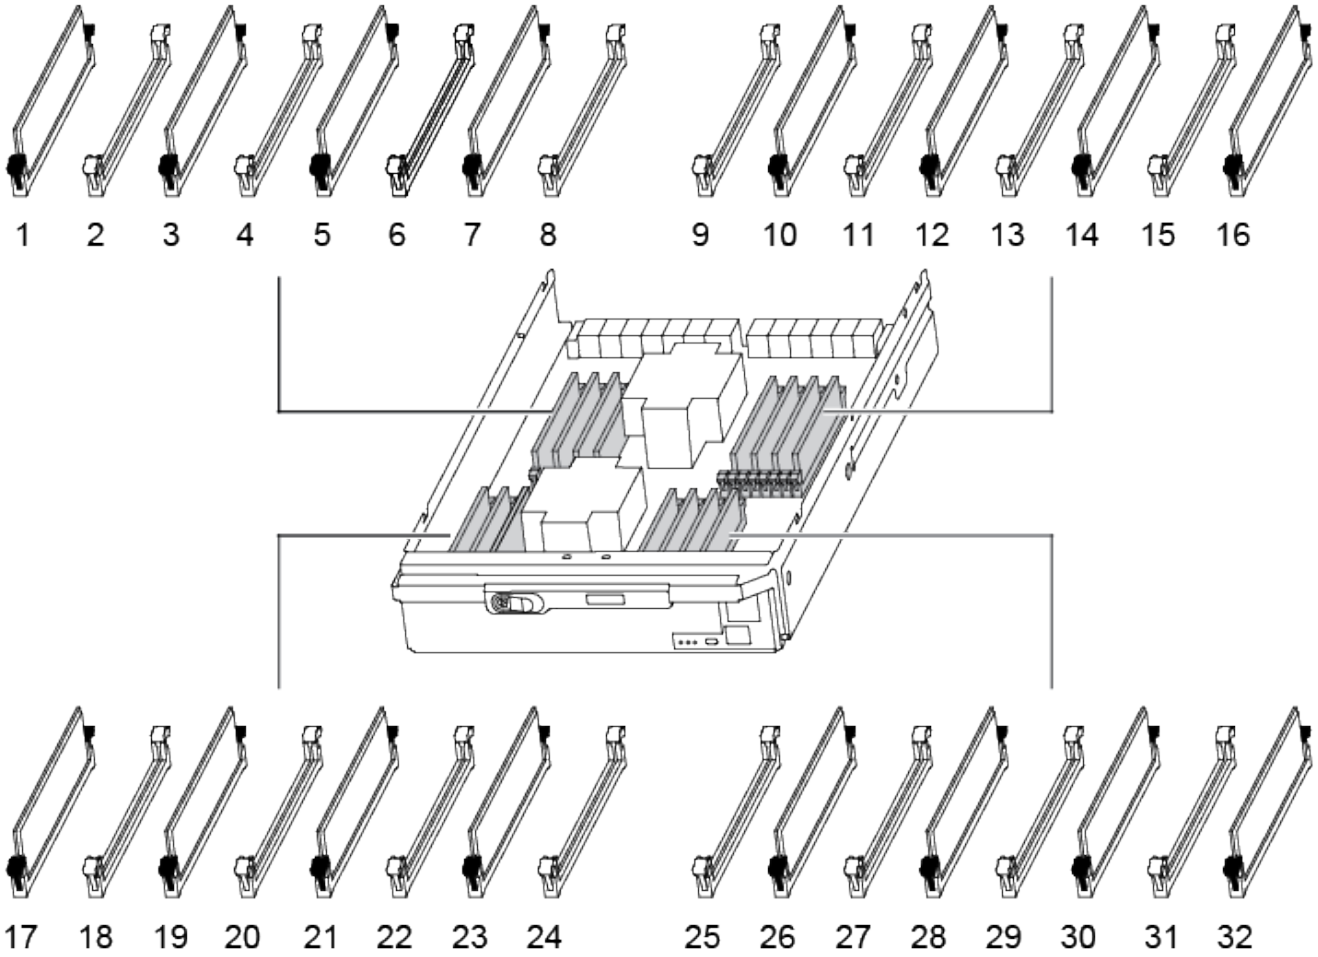

3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

Animation - Replace DIMM

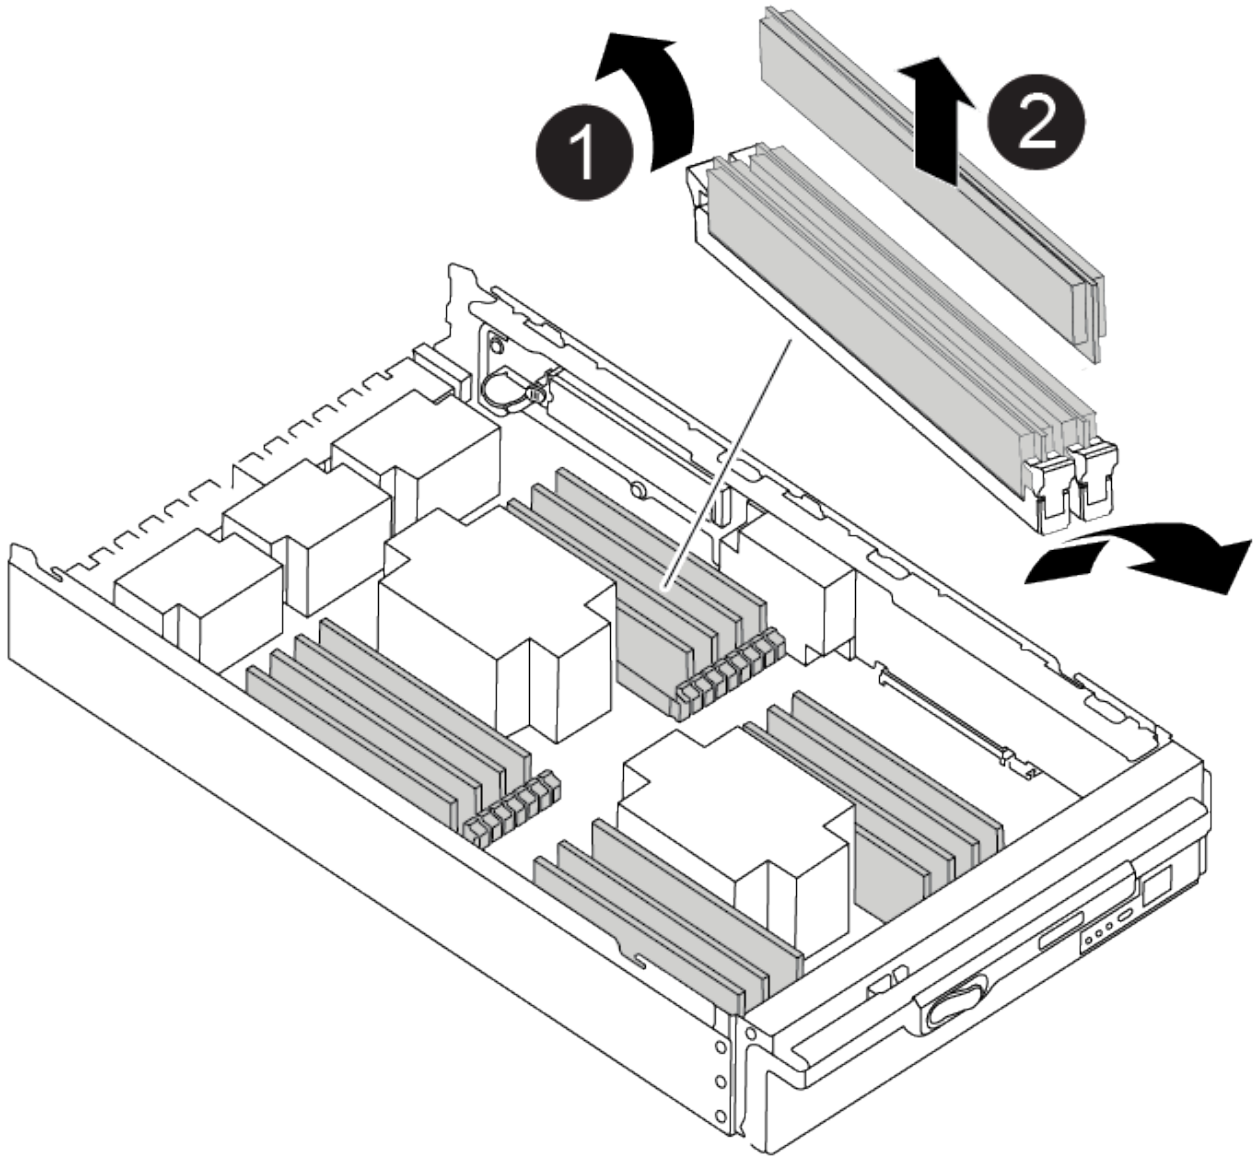

1	DIMM ejector tabs
2	DIMM

4. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

5. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Close the controller module cover.

Step 4: Install the controller

After you install the components into the controller module, you must install the controller module back into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.

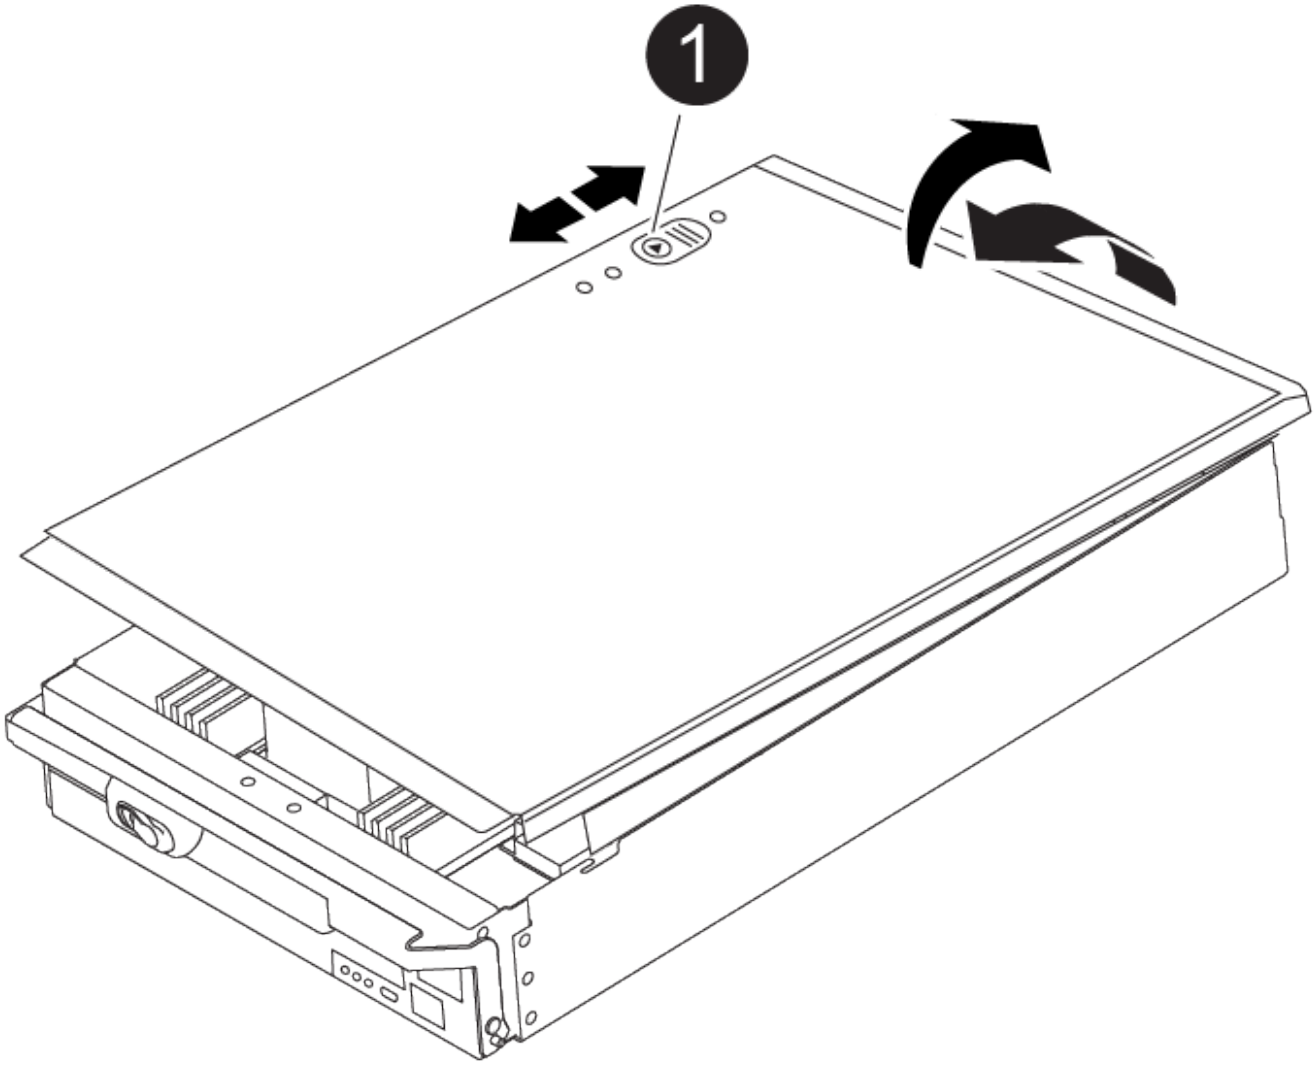

1	Controller module cover locking button

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Animation - Install controller

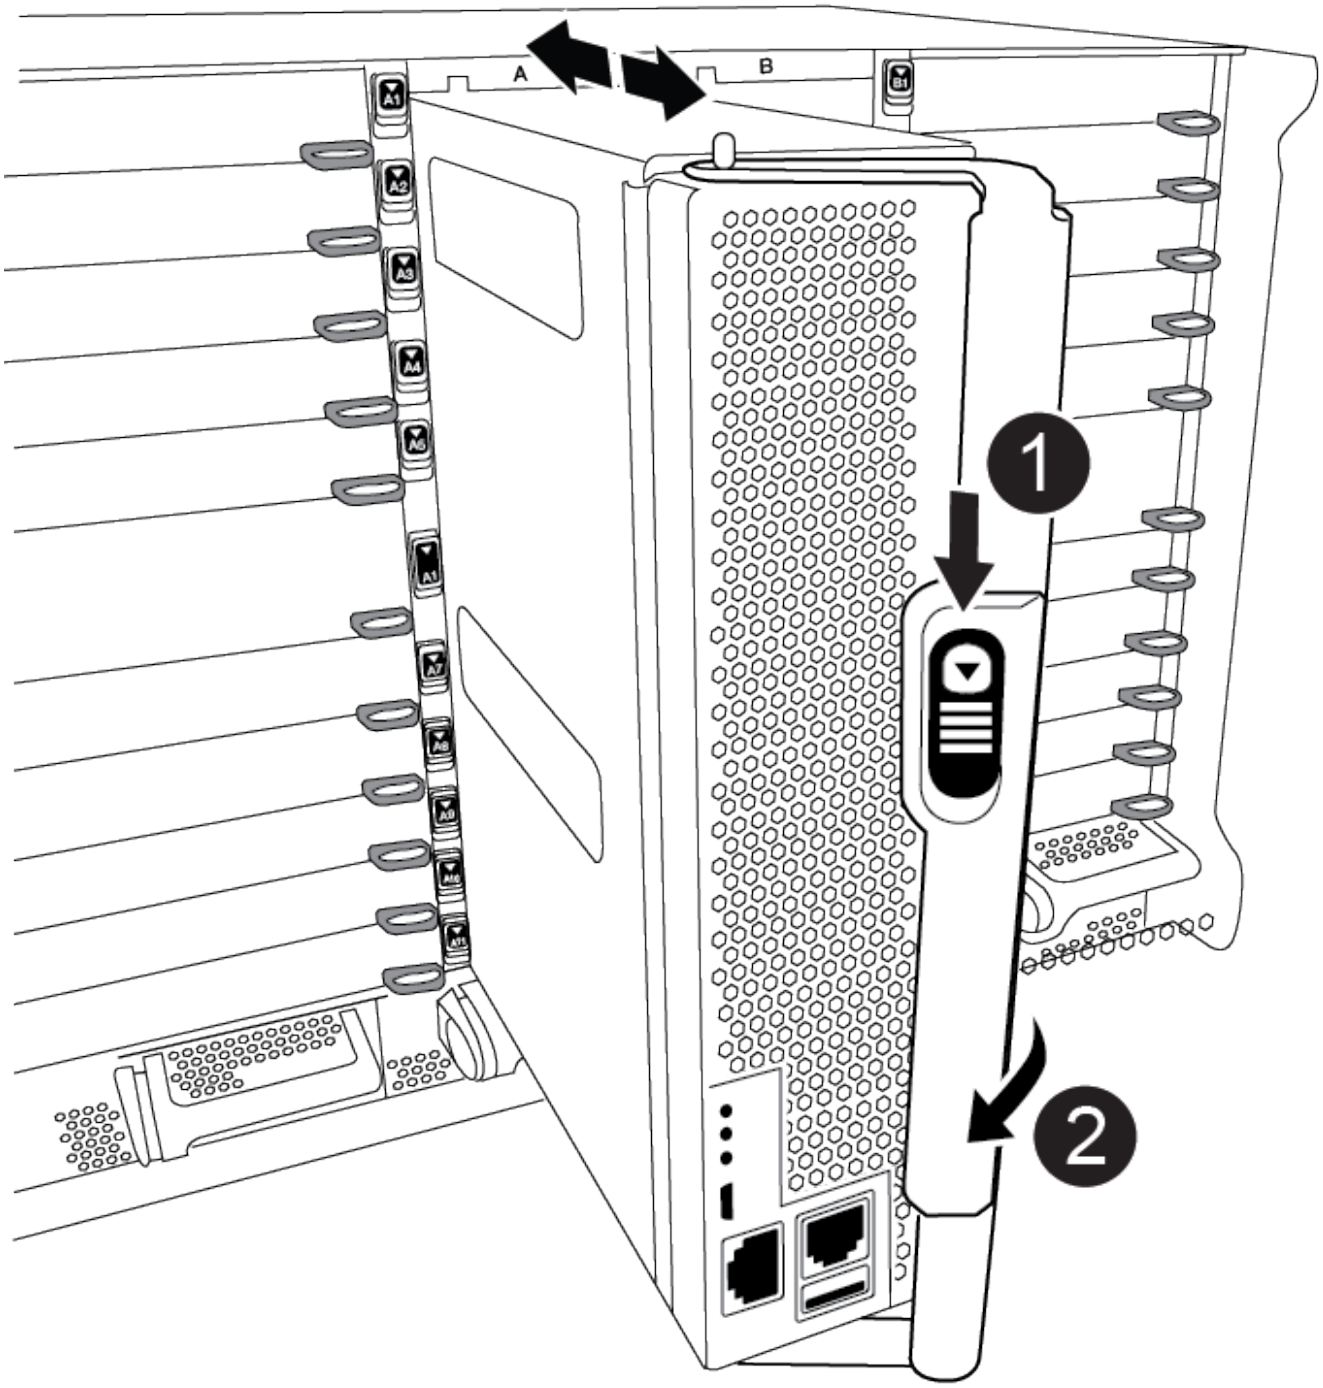

1	Cam handle release button
2	Cam handle

Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   d. Interrupt the boot process by pressing `Ctrl-C` when you see Press Ctrl-C for Boot Menu.
   e. Select the option to boot to Maintenance mode from the displayed menu.

Step 5: Run system-level diagnostics

After installing a new DIMM, you should run diagnostics.

Your system must be at the LOADER prompt to start System Level Diagnostics.

All commands in the diagnostic procedures are issued from the controller where the component is being replaced.

1. If the controller to be serviced is not at the LOADER prompt, perform the following steps:
   a. Select the Maintenance mode option from the displayed menu.
   b. After the controller boots to Maintenance mode, halt the controller: `halt`

      After you issue the command, you should wait until the system stops at the LOADER prompt.

      During the boot process, you can safely respond `y` to prompts.

      - If a prompt appears warning that when entering Maintenance mode in an HA configuration, you must ensure that the healthy controller remains down.

2. At the LOADER prompt, access the special drivers specifically designed for system-level diagnostics to function properly: `boot_diags`

   During the boot process, you can safely respond `y` to the prompts until the Maintenance mode prompt (*>) appears.

3. Run diagnostics on the system memory: `sldiag device run -dev mem`
4. Verify that no hardware problems resulted from the replacement of the DIMMs: `sldiag device status`

`-dev mem -long -state failed`

System-level diagnostics returns you to the prompt if there are no test failures, or lists the full status of failures resulting from testing the component.

5. Proceed based on the result of the preceding step:

If the system-level diagnostics tests…	Then…
Were completed without any failures	a. Clear the status logs: `sldiag device clearstatus` b. Verify that the log was cleared: `sldiag device status` The following default response is displayed: SLDIAG: No log messages are present. c. Exit Maintenance mode: `halt` The controller displays the LOADER prompt. d. Boot the controller from the LOADER prompt: `bye` e. Return the controller to normal operation:

If your controller is in…	Then…
An HA pair	Perform a give back: `storage failover giveback -ofnode replacement_node_name` **Note:** If you disabled automatic giveback, re-enable it with the storage failover modify command.

If your controller is in…	Then…
Resulted in some test failures	Determine the cause of the problem: a. Exit Maintenance mode: `halt` After you issue the command, wait until the system stops at the LOADER prompt. b. Verify that you have observed all the considerations identified for running system-level diagnostics, that cables are securely connected, and that hardware components are properly installed in the storage system. c. Boot the controller module you are servicing, interrupting the boot by pressing `Ctrl-C` when prompted to get to the Boot menu: ◦ If you have two controller modules in the chassis, fully seat the controller module you are servicing in the chassis. The controller module boots up when fully seated. ◦ If you have one controller module in the chassis, connect the power supplies, and then turn them on. d. Select Boot to maintenance mode from the menu. e. Exit Maintenance mode by entering the following command: `halt` After you issue the command, wait until the system stops at the LOADER prompt. f. Rerun the system-level diagnostic test.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the Destage Control Power Module containing the NVRAM11 battery - AFF A900

To hot-swap a destage controller power module (DCPM), which contains the NVRAM11 battery, you must locate the failed DCPM module, remove it from the chassis, and install the replacement DCPM module.

You must have a replacement DCPM module in-hand before removing the failed module from the chassis and it must be replaced within five minutes of removal. Once the DCPM module is removed from the chassis, there is no shutdown protection for the controller module that owns the DCPM module, other than failover to the other controller module.

Step 1: Replace the DCPM module

To replace the DCPM module in your system, you must remove the failed DCPM module from the system and then replace it with a new DCPM module.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel on the front of the system and set it aside.
3. Locate the failed DCPM module in the front of the system by looking for the Attention LED on the module.

   The LED will be steady amber if the module is faulty.

   > The DCPM module must be replaced in the chassis within five minutes of removal or the associated controller will shut down.

4. Press the terra cotta release button on the module handle, and then slide the DCPM module out of the chassis.

   Animation - Remove/install DCPM

   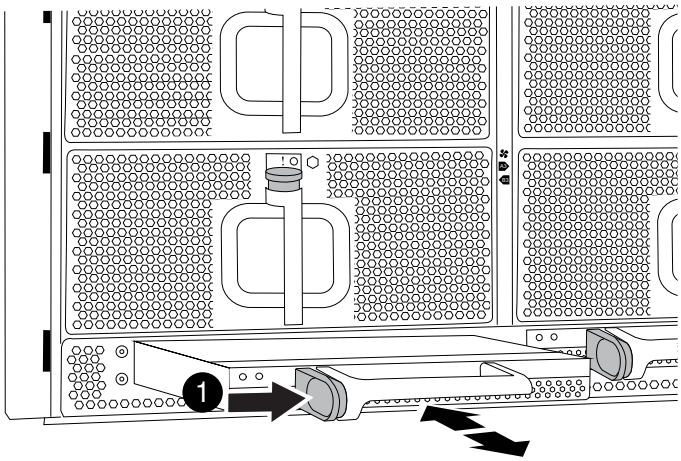

1	DCPM module terra cotta locking button

5. Align the end of the DCPM module with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   > The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

   The Amber LED flashes four times upon insertion and the green LED also flashes if the battery is providing a voltage. If it does not flash, it will likely need to be replaced.

Step 2: Dispose of batteries

You must dispose of batteries according to the local regulations regarding battery recycling or disposal. If you cannot properly dispose of batteries, you must return the batteries to NetApp, as described in the RMA instructions that are shipped with the kit.

Safety Information and Regulatory Notices

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a fan - AFF A900

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press the terra cotta button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

> The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

Animation - Remove/install fan

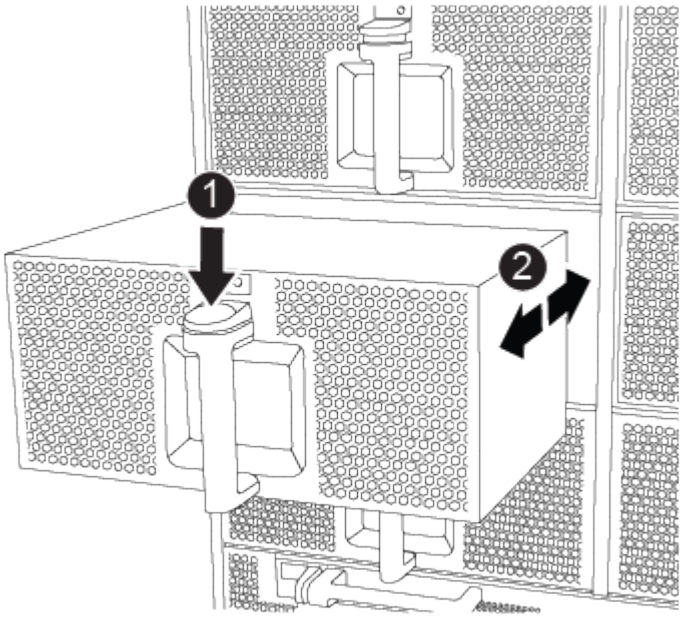

1	Terra cotta release button

2	Slide fan in/out of chassis

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Add an I/O module - AFF A900

You can add an I/O module to your AFF A900 storage system when there are empty slots available or when all slots are fully populated.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 2: MetroCluster configuration

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster

Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Add the new I/O module

If the storage system has empty slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- To non-disruptively add an I/O module, you must take over the target controller, remove the slot blanking cover in the target slot or remove an existing I/O module, add the new or replacement I/O module, and then giveback the target controller.
- Make sure that all other components are functioning properly.

Add I/O module to an empty slot

You can add a new I/O module into a storage system with available empty slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the target slot blanking cover:
   a. Depress the lettered and numbered cam latch.
   b. Rotate the cam latch down until it is the open position.
   c. Remove the blanking cover.
3. Install the I/O module:
   a. Align the I/O module with the edges of the slot.
   b. Slide the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
4. If the replacement I/O module is a NIC, cable the module to the data switches.

   Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

5. Reboot the controller from the LOADER prompt: *bye*

   This reinitializes the PCIe cards and other components and reboots the node.

6. Give back the controller from the partner controller. `storage failover giveback -ofnode target_node_name`
7. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto -giveback true`
8. If you are using slots 3 and/or 7 for networking, use the `storage port modify -node <node name> -port <port name> -mode network` command to convert the slot for networking use.
9. Repeat these steps for controller B.
10. If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam latch.

      The cam latch moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

      Animation - Remove or replacing an I/O module

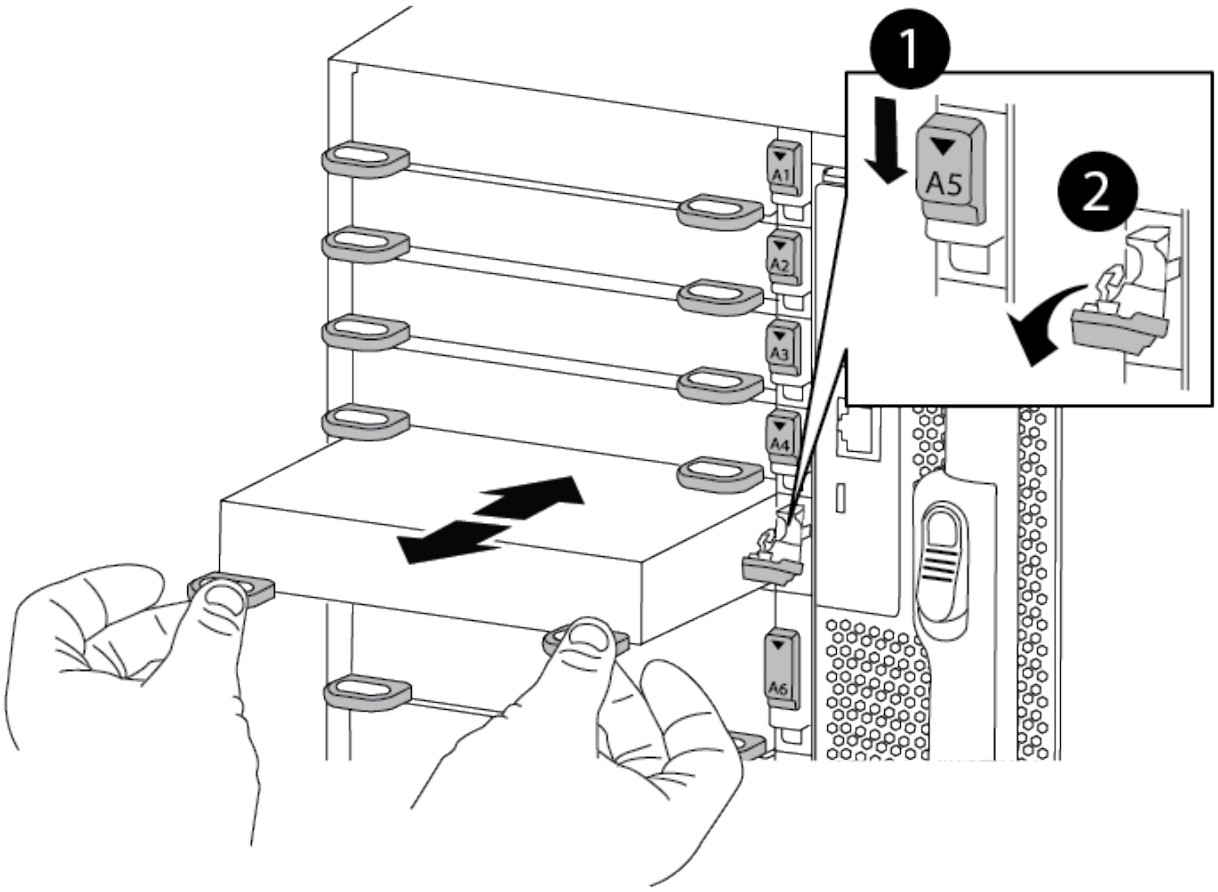

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Install the I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Slide the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
5. Repeat the remove and install steps to replace additional modules for controller A.
6. If the replacement I/O module is a NIC, cable the module or modules to the data switches.
7. Reboot the controller from the LOADER prompt:
   a. Check the version of BMC on the controller: `system service-processor show`
   b. Update the BMC firmware if needed: `system service-processor image update`
   c. Reboot the node: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.

If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement

8. Give back the controller from the partner controller. `storage failover giveback -ofnode target_node_name`
9. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto-giveback true`
10. If you added:

If I/O module is a…	Then…
NIC module in slots 3 or 7,	Use the `storage port modify -node *<node name> -port *<port name> -mode network` command for each port.
Storage module	Install and cable your NS224 shelves, as described in Hot-add workflow.

11. Repeat these steps for controller B.

Replace an I/O module - AFF A900

To replace an I/O module, you must perform a specific sequence of tasks.

- You can use this procedure with all versions of ONTAP supported by your system.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired node

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster

Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace I/O modules

To replace an I/O module, locate it within the chassis and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

3. Remove the target I/O module from the chassis:

   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

Make sure that you keep track of which slot the I/O module was in.

Animation - Remove/install I/O module

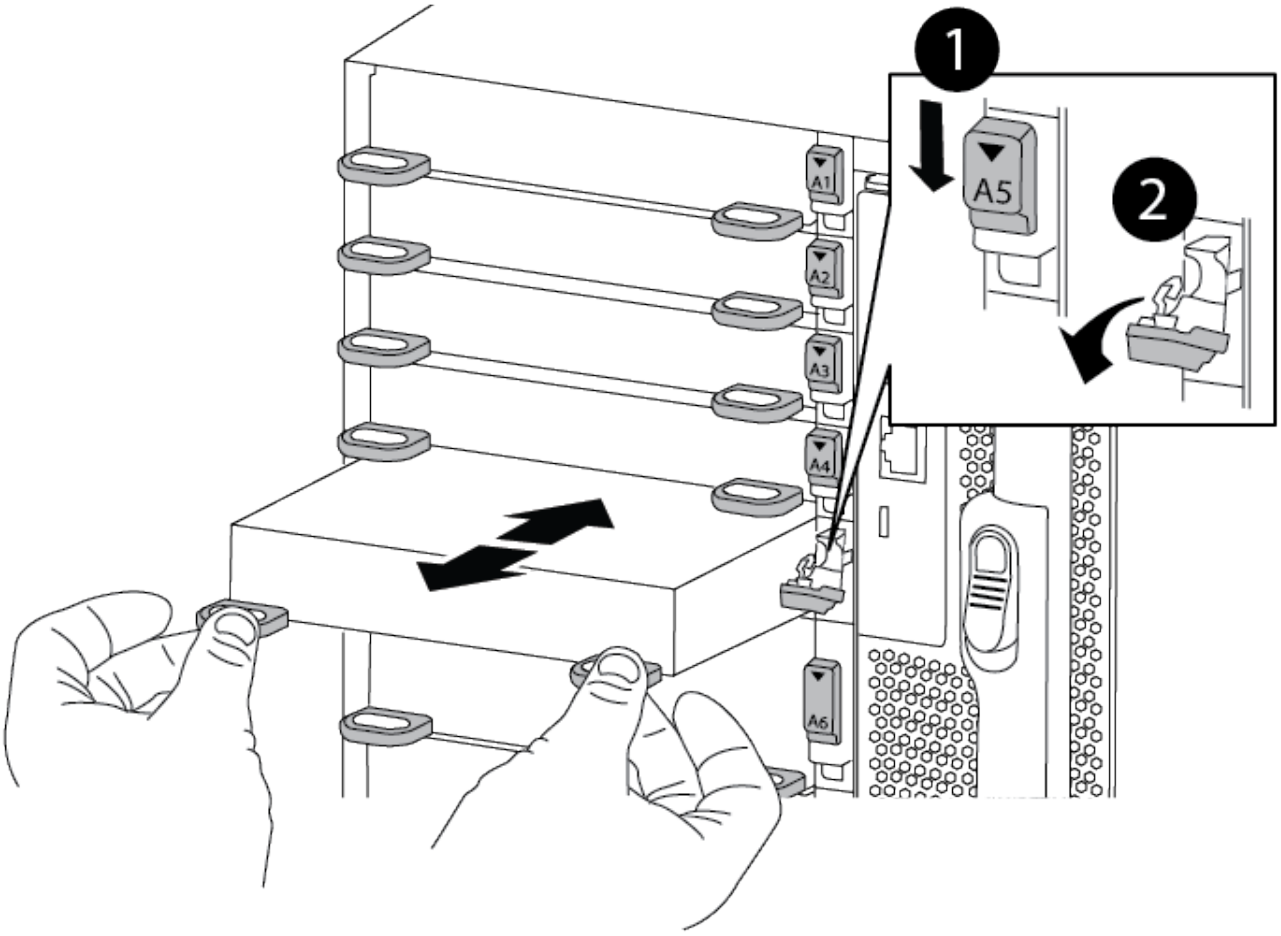

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Install the replacement I/O module into the chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.
6. Recable the I/O module, as needed.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller module.

If the new I/O module is not the same model as the failed module, you must first reboot the BMC.

Steps

1. Reboot the BMC if the replacement module is not the same model as the old module:
   a. From the LOADER prompt, change to advanced privilege mode: `priv set advanced`
   b. Reboot the BMC: `sp reboot`
2. From the LOADER prompt, reboot the node: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.
3. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs, convert these ports to 10 GbE connections by using the nicadmin convert command from Maintenance mode. See Convert 40GbE NIC ports into multiple 10GbE ports for 10GbE connectivity for more information.

   Be sure to exit Maintenance mode after completing the conversion.
4. Return the node to normal operation: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an LED USB module - AFF A900

The LED USB module provides connectivity to console ports and system status. Replacement of this module does not require tools and does not interrupt service.

Step 1: Replace the LED USB module

Steps

1. Remove the impaired LED USB module:

   Animation - Remove/install LED-USB module

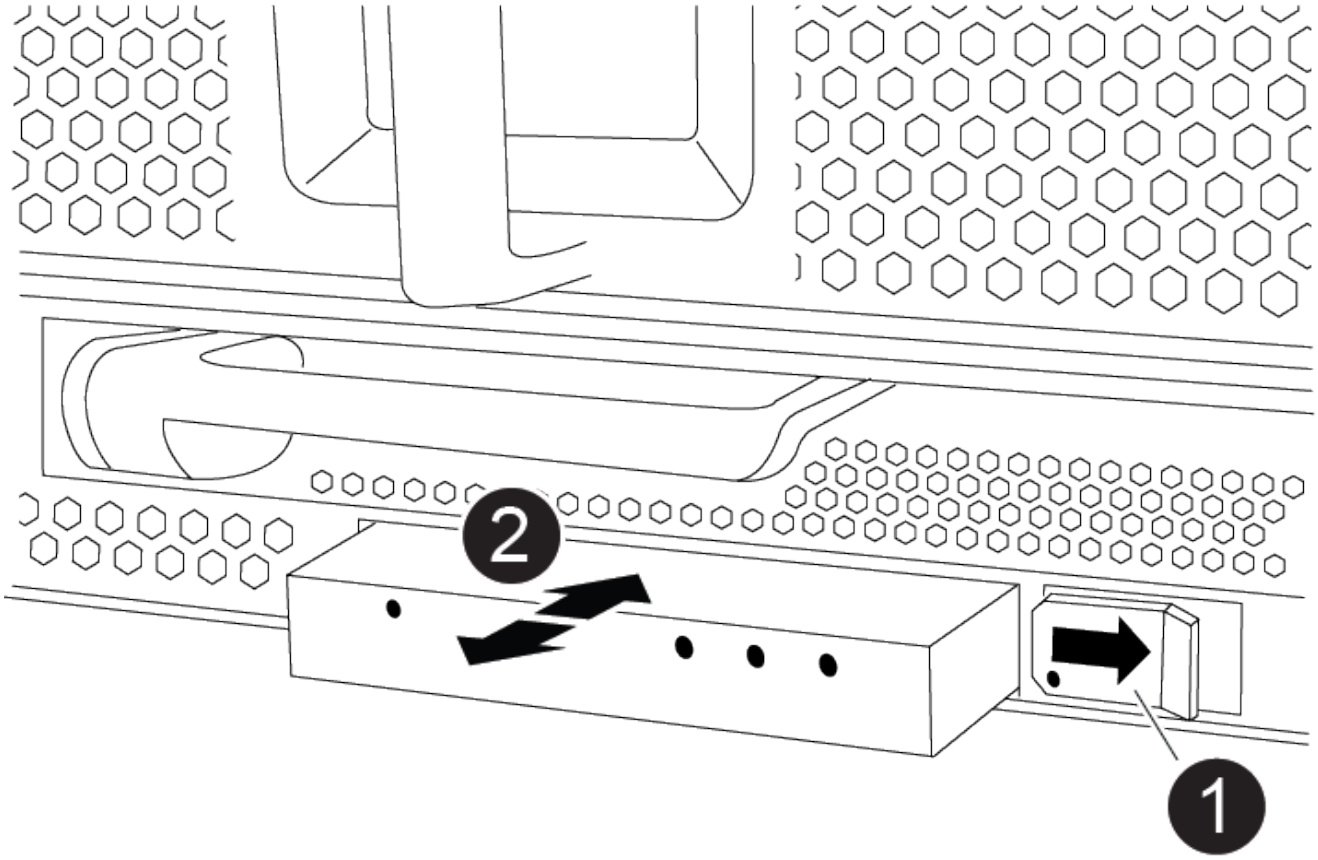

1	Locking button
2	USB LED module

a. With the bezel removed, locate the LED USB module at the front of the chassis, on the bottom left side.

b. Slide the latch to partially eject the module.

c. Pull the module out of the bay to disconnect it from the midplane. Do not leave the slot empty.

2. Install the new LED USB module:

a. Align the module to the bay with the notch in the corner of the module positioned near the slider latch on the chassis. The bay will prevent you from installing the module upside down.

b. Push the module into the bay until it is fully seated flush with the chassis.

There is an audible click when the module is secure and connected to the midplane.

Step 2: Return the failed component

1. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVRAM module and/or NVRAM DIMMs - AFF A900

The NVRAM module consists of the NVRAM11 and DIMMs. You can replace a failed

NVRAM module or the DIMMs inside the NVRAM module. To replace a failed NVRAM module, you must remove it from the chassis, move the DIMMs to the replacement module, and install the replacement NVRAM module into the chassis.

To replace and NVRAM DIMM, you must remove the NVRAM module from the chassis, replace the failed DIMM in the module, and then reinstall the NVRAM module.

About this task

Because the system ID is derived from the NVRAM module, if replacing the module, disks belonging to the system are reassigned to a new system ID.

Before you begin

- All disk shelves must be working properly.
- If your system is in an HA pair, the partner controller must be able to take over the controller associated with the NVRAM module that is being replaced.
- This procedure uses the following terminology:
  - The impaired controller is the controller on which you are performing maintenance.
  - The healthy controller is the HA partner of the impaired controller.
- This procedure includes steps for automatically reassigning disks to the controller module associated with the new NVRAM module. You must reassign the disks when directed to in the procedure. Completing the disk reassignment before giveback can cause issues.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You cannot change any disks or disk shelves as part of this procedure.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 6 in the chassis and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the target NVRAM module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The NVRAM module disengages from the chassis and moves out a few inches.

   c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

      Animation - Replace the NVRAM module

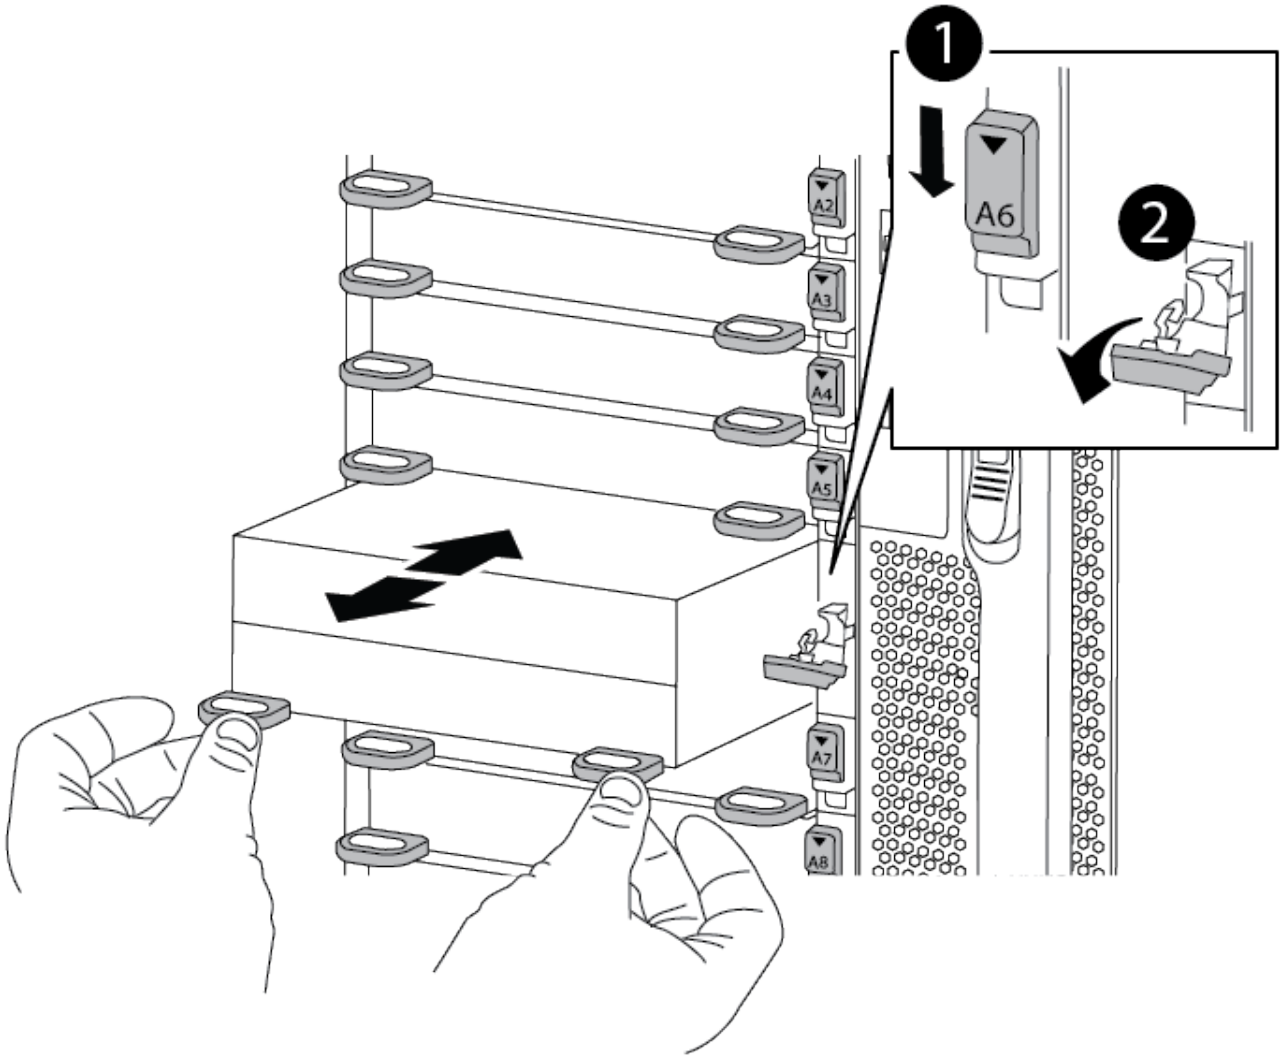

1	Lettered and numbered cam latch
2	Cam latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

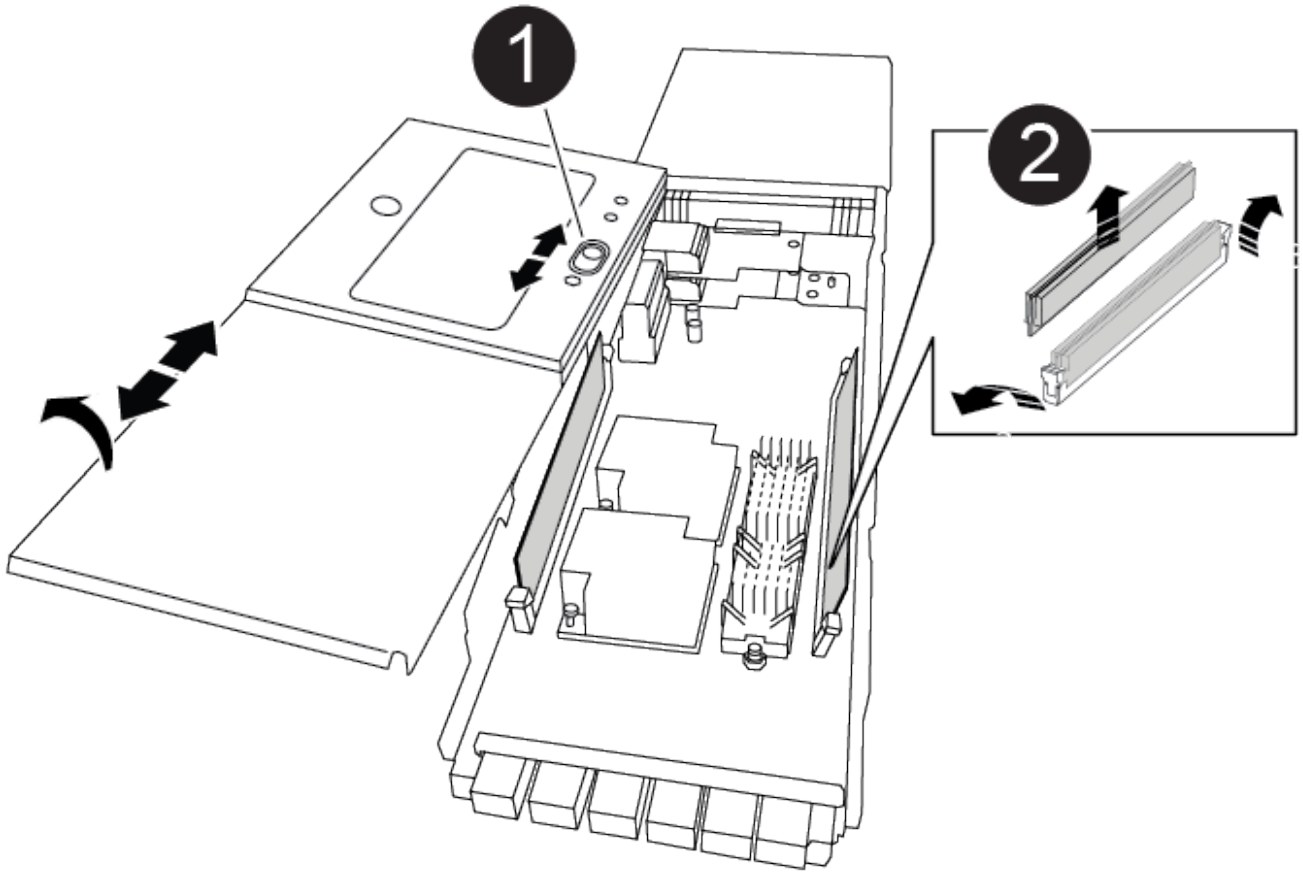

1	Cover locking button
2	DIMM and DIMM ejector tabs

4. Remove the DIMMs, one at a time, from the old NVRAM module and install them in the replacement NVRAM module.
5. Close the cover on the module.
6. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered cam latch begins to engage with the I/O cam pin, and then push the cam latch all the way up to lock the module in place.

Step 3: Replace a NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, you must remove the NVRAM module, open the module, and then replace the target DIMM.

1. If you are not already grounded, properly ground yourself.
2. Remove the target NVRAM module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

b. Rotate the cam latch down until it is in a horizontal position.

   The NVRAM module disengages from the chassis and moves out a few inches.

c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

   Animation - Replace NVRAM DIMM

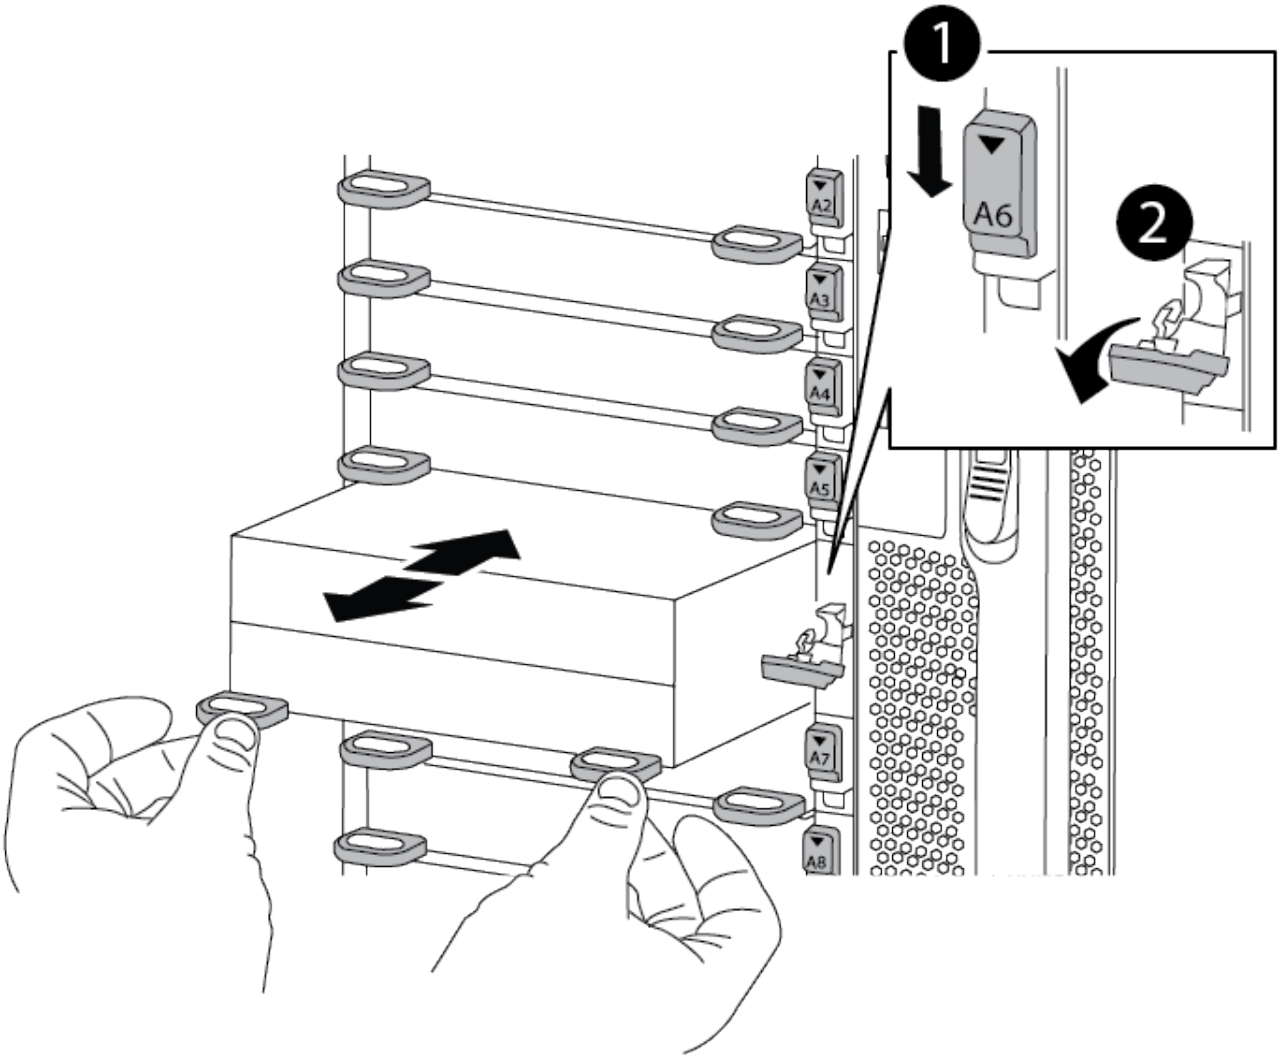

1	Lettered and numbered cam latch
2	cam latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

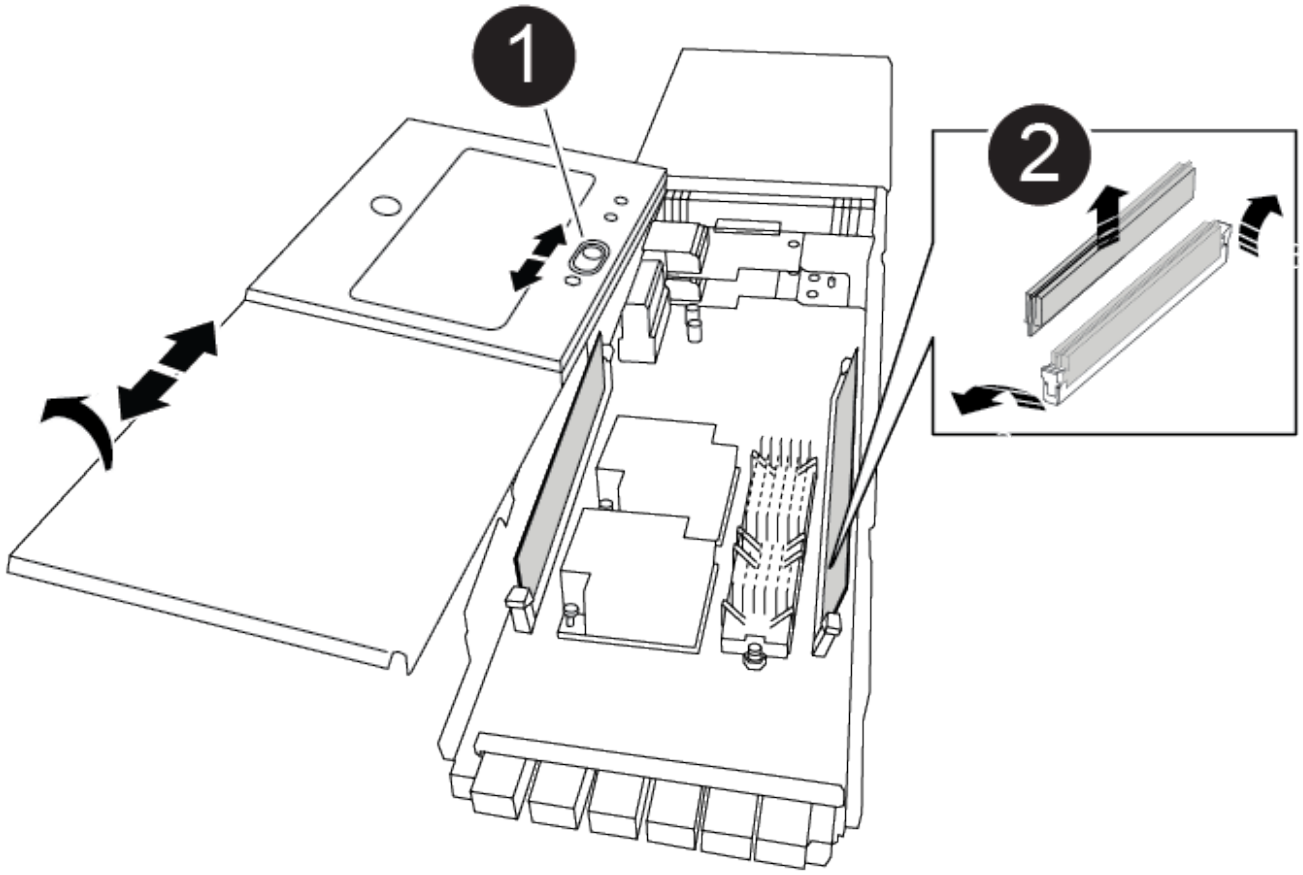

1	Cover locking button
2	DIMM and DIMM ejector tabs

4. Locate the DIMM to be replaced inside the NVRAM module, and then remove it by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
5. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
6. Close the cover on the module.
7. Install the NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered cam latch begins to engage with the I/O cam pin, and then push the cam latch all the way up to lock the module in place.

Step 4: Reboot the controller

After you replace the FRU, you must reboot the controller module.

1. To boot ONTAP from the LOADER prompt, enter `bye`.

Step 5: Reassign disks

You must confirm the system ID change when you boot the replacement controller and then verify that the change was implemented.

> Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the replacement controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the replacement controller, boot the controller and entering y if you are prompted to override the system ID due to a system ID mismatch.
3. Wait until the Waiting for giveback… message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1:> storage failover show
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The replacement controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

5. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the replacement controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 151759706:

```
node1:> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        151759706  151759706  -
151759706 Pool0
1.0.1  aggr0_1  node1 node1           151759706  151759706  -
151759706 Pool0
.
.
.
```

6. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

7. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The replacement controller is the current owner of the disks on the disaster site.

     See Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration for more information.

8. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

9. Verify that the expected volumes are present for each controller: `vol show -node node-name`
10. If storage encryption is enabled, you must restore functionality.
11. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - AFF A900

Swapping out a power supply involves turning off, disconnecting, and removing the power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

About this task

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- There are four power supplies in the system.
- Power supplies are auto-ranging.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.

3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
4. Press and hold the terra cotta button on the power supply handle, and then pull the power supply out of the chassis.

   CAUTION:

   When removing a power supply, always use two hands to support its weight.

   Animation - Remove/install PSU

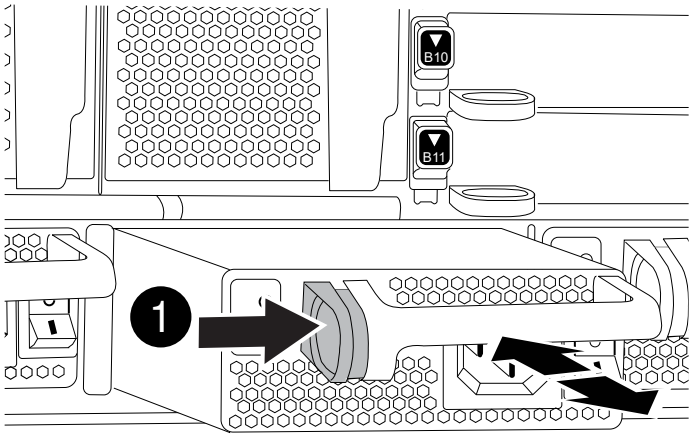

1	Locking button

5. Make sure that the on/off switch of the new power supply is in the Off position.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply.
   b. Secure the power cable to the power supply using the power cable retainer.

      Once power is restored to the power supply, the status LED should be green.

8. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

The green power LED lights when the PSU is fully inserted into the chassis and the amber attention LED flashes initially, but turns off after a few moments.

9. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replacing the real-time clock battery - AFF A900

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove the controller

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

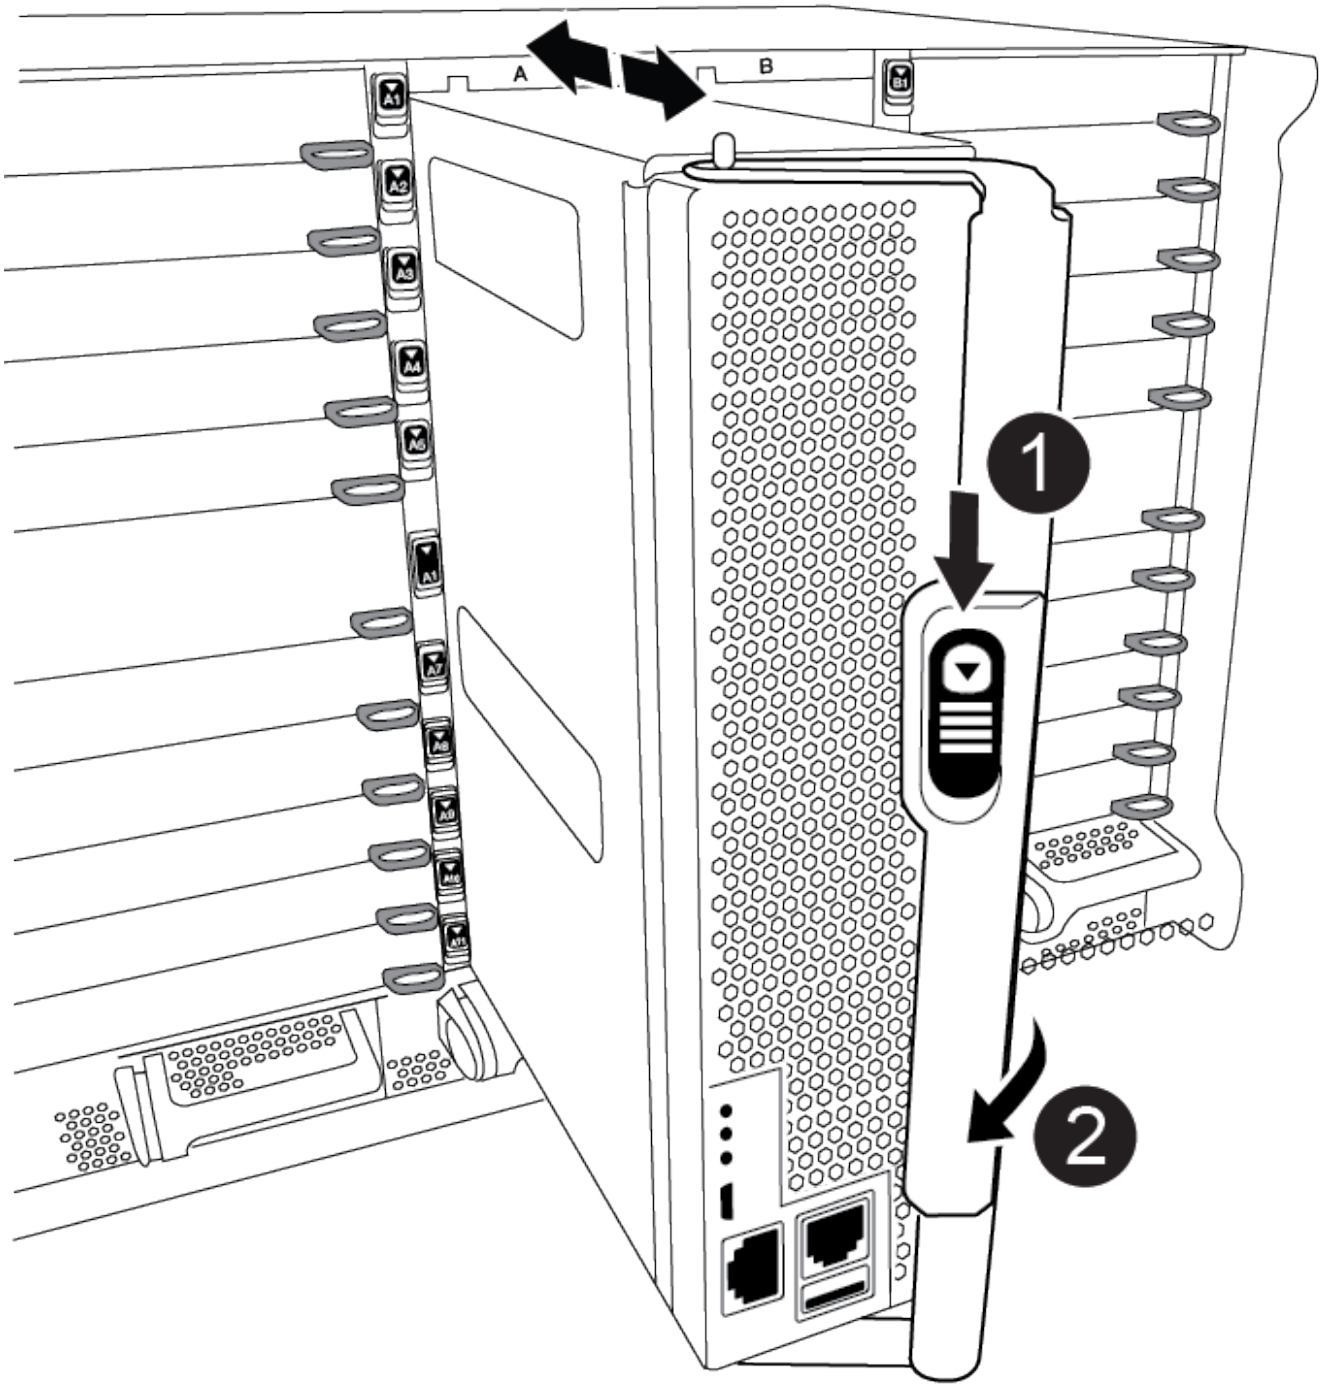

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

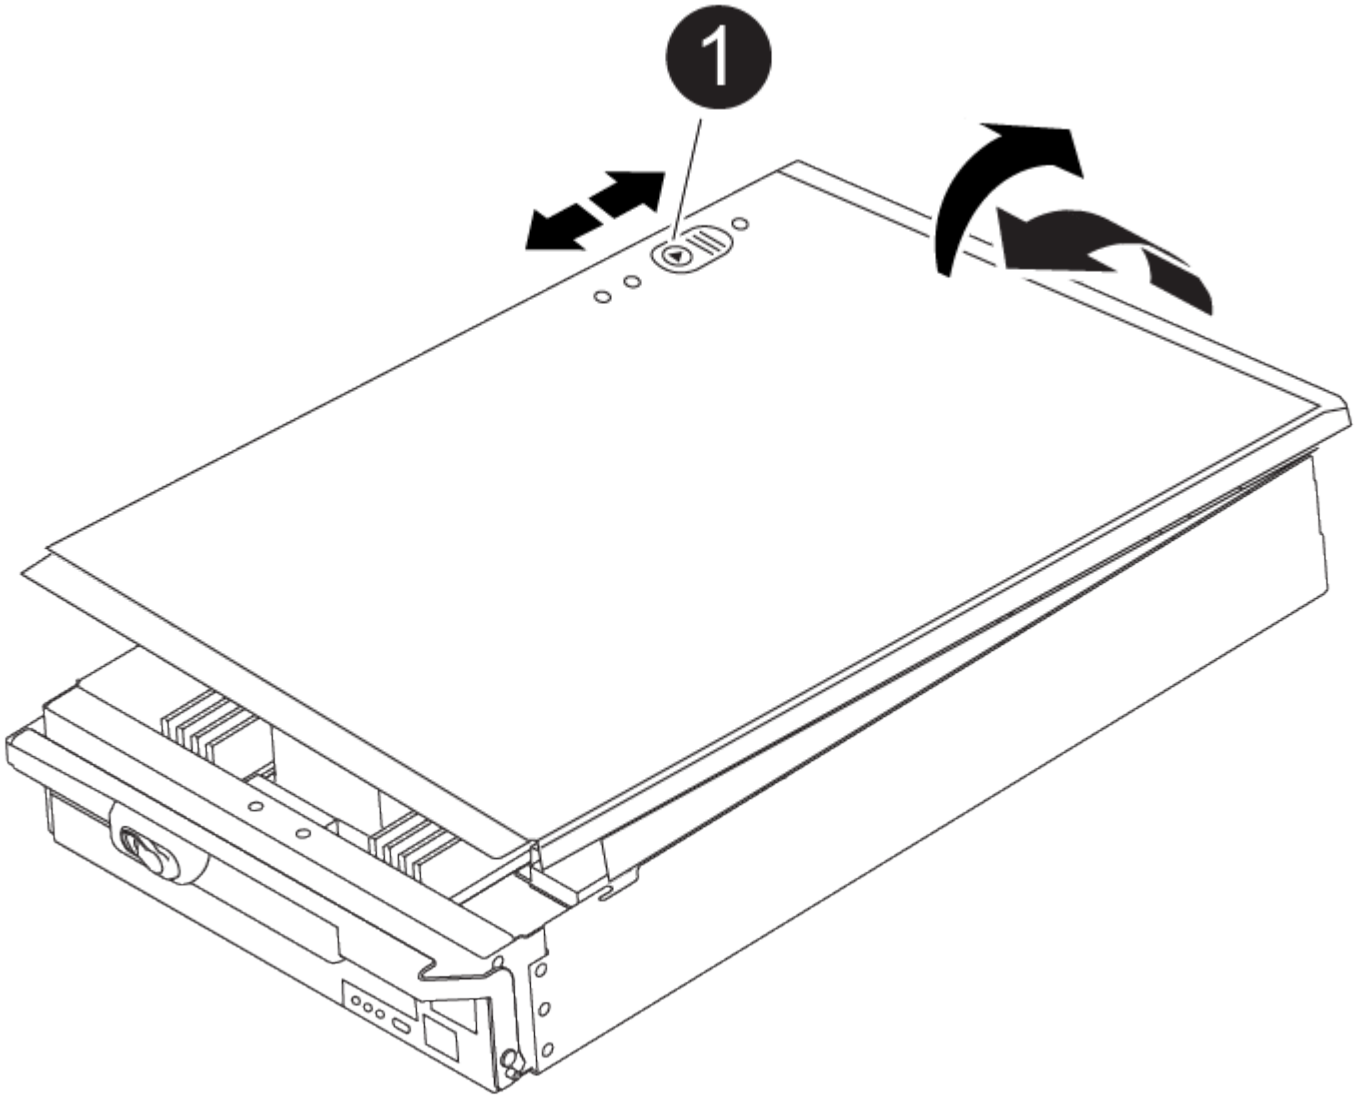

1	Controller module cover locking button

Step 3: Replace the RTC battery

To replace the RTC battery, you must locate the failed battery in the controller module, remove it from the holder, and then install the replacement battery in the holder.

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

   Animation - Replace RTC battery

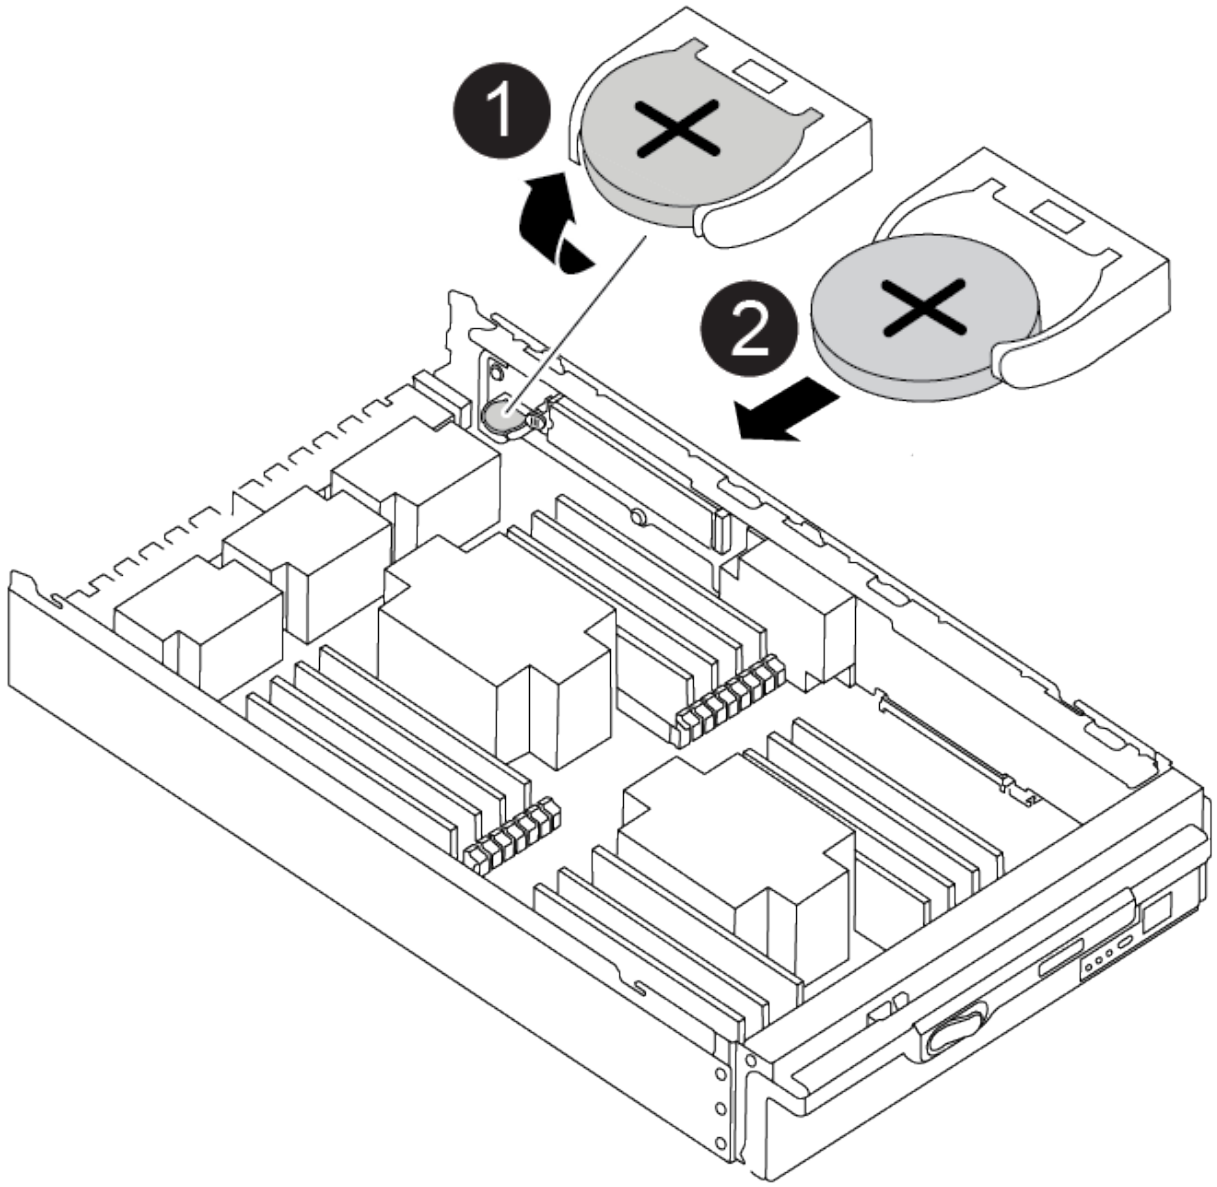

1	RTC battery
2	RTC battery housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.

7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
8. Reinstall the controller module cover.

Step 4: Reinstall the controller module and set time/date

After you replace the RTC battery, you must reinstall the controller module. If the RTC battery has been left out of the controller module for more than 10 minutes, you may have to reset the time and date.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.

If your system stops at the boot menu, select the option for "Reboot node" and respond y when prompted,
then boot to LOADER by pressing `Ctrl-C`.

1. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
2. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

3. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF C-Series systems

AFF C30 and AFF C60 systems

Install and setup

Installation and setup workflow - AFF C30 and AFF C60

To install and set up your AFF C30 or AFF C60 storage system, you must review the installation requirements, prepare your site, install and cable the hardware components, power on the storage system, and set up the ONTAP cluster.

Review the installation requirements

Before installing your storage system, it must meet the installation requirements.

Prepare for installation

To prepare for installation, get the site ready, check environmental and electrical requirements, and ensure there's enough rack space. Then, unpack the equipment, compare contents to the packing slip, and register the hardware to access support benefits.

Install the hardware

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system and shelves in the cabinet or telco rack.

Cable the hardware

To cable the hardware, connect the controllers to your network and then to your shelves.

Power on the storage system

To power on your storage system, power on each shelf and assign a unique shelf ID as needed, and then power on the controllers.

6 Set up your cluster

After you've powered on your storage system, you set up your cluster.

Installation requirements - AFF C30 and AFF C60

Review the requirements for your AFF C30 or AFF C60 storage system.

Equipment needed for install

To install your storage system, you need the following equipment and tools.

- Access to a Web browser to configure your storage system
- Electrostatic discharge (ESD) strap
- Flashlight
- Laptop or console with a USB/serial connection
- Phillips #2 screwdriver

Lifting precautions

Storage systems and shelves are heavy. Exercise caution when lifting and moving these items.

Storage system weight

Take the necessary precautions when moving or lifting your storage system.

An A1K storage system can weigh up to 62.83 lbs (28.5 kg). To lift the storage system, use two people or a hydraulic lift.

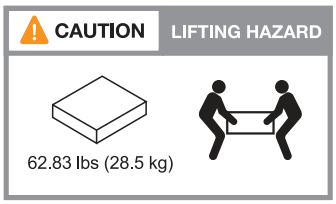

The storage system can weigh up to 61.5 lbs (27.9 kg). To lift the storage system, use two people or a hydraulic lift.

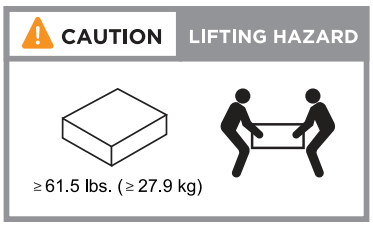

Shelf weight

Take the necessary precautions when moving or lifting your shelf.

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

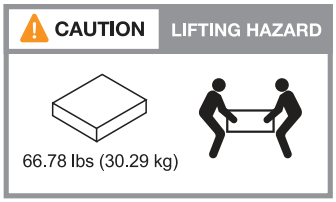

An NS224 shelf with NSM100B modules can weigh up to 56.8 lbs (25.8 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

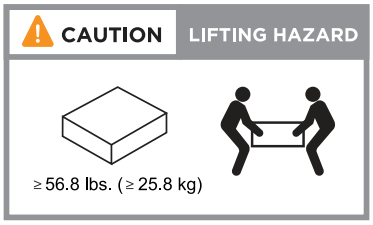

Related information

- Safety information and regulatory notices

What's next?

After you've reviewed the installation requirements and considerations for your storage system, you prepare for installation.

Prepare to install - AFF C30 and AFF C60

Prepare to install your AFF C30 or AFF C60 storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the storage system to access support benefits.

Step 1: Prepare the site

To install your storage system, ensure that the site and the cabinet or rack that you plan to use meet specifications for your configuration.

Steps

1. Use NetApp Hardware Universe to confirm that your site meets the environmental and electrical requirements for your storage system.
2. Make sure you have adequate cabinet or rack space for your storage system, shelves, and any switches:
   - 4U in an HA configuration
   - 2U for each NS224 storage shelf
   - 2U for a storage system
   - 2U for each NS224 storage shelf
   - 1U for most switches
3. Install any required network switches.

   See the Switch documentation for installation instructions and NetApp Hardware Universe for compatibility information.

Step 2: Unpack the boxes

After you've ensured that the site and the cabinet or rack that you plan to use for your storage system meet the required specifications, unpack all boxes and compare the contents to the items on the packing slip.

Steps

1. Carefully open all the boxes and lay out the contents in an organized manner.
2. Compare the contents you've unpacked with the list on the packing slip.

   You can get your packing list by scanning the QR code on the side of the shipping carton.

   The following items are some of the contents you might see in the boxes.

   Ensure that everything in the boxes matches the list on the packing slip. If there are any discrepancies, note them down for further action.

Hardware

- Bezel
- Cable management device
- Storage system
- Rail kits with instructions (optional)
- Storage shelf (if you ordered additional storage)

Cables

- Management Ethernet cables (RJ-45 cables)
- Network cables
- Power cords
- Storage cables (if you ordered additional storage)
- USB-C serial console cable

Step 3: Register your storage system

After you've ensured that your site meets the requirements for your storage system specifications, and you've verified that you have all the parts you ordered, you should register your storage system.

Steps

1. Locate the System Serial Numbers (SSN) for every controller being installed. You can find the serial numbers in the following locations:
2. You can find the serial numbers in the following locations:
   - On the packing slip
   - In your confirmation email
   - On each controller's System Management module
   - On each controller

3. Go to the NetApp Support Site.
4. Determine whether you need to register your storage system:

If you are a…	Follow these steps…
Existing NetApp customer	a. Sign in with your username and password. b. Select **Systems** > **My Systems**. c. Confirm that the new serial numbers are listed. d. If it is not, follow the instructions for new NetApp customers.
New NetApp customer	a. Click **Register Now**, and create an account. b. Select **Systems** > **Register Systems**. c. Enter the storage system's serial numbers and requested details. After your registration is approved, you can download any required software. The approval process might take up to 24 hours.

What's next?

After you've prepared to install your storage system, you install the hardware for your storage system.

Install the hardware - AFF C30 and AFF C60

After you prepare to install your AFF C30 or AFF C60 storage system, install the hardware for the storage system. First, install the rail kits. Then install and secure your storage system in a cabinet or telco rack.

Skip this step if your storage system came in a cabinet.

Before you begin

- Make sure you have the instructions packaged with the rail kit.
- Be aware of the safety concerns associated with the weight of the storage system and shelf.
- Understand that the airflow through the storage system enters from the front where the bezel or end caps are installed and exhausts out the rear where the ports are located.

Steps

1. Install the rail kits for your storage system and shelves as needed, using the instructions included with the kits.
2. Install and secure your storage system in the cabinet or telco rack:
   a. Position the storage system onto the rails in the middle of the cabinet or telco rack, and then support the storage system from the bottom and slide it into place.
   b. Make sure that the guiding pins of the cabinet or telco rack are securley in the chassis guide slots.
   c. Secure the storage system to the cabinet or telco rack using the included mounting screws.
3. Attach the bezel to the front of the storage system.
4. Attach the cable management devices to the rear of the storage system.

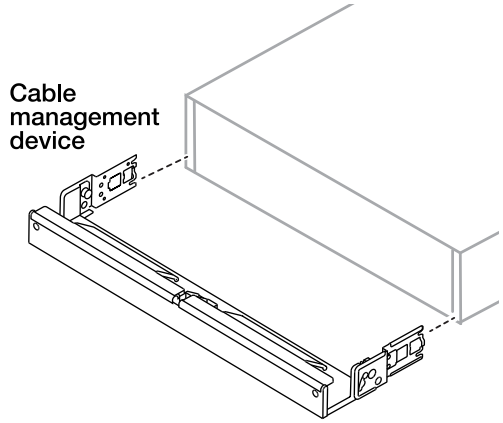

5. Install and secure the shelf as needed.
   a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

      If you are installing multiple shelves, place the first shelf directly above the controllers. Place the second shelf directly under the controllers. Repeat this pattern for any additional shelves.

   b. Secure the shelf to the cabinet or telco rack using the included mounting screws.

      What's next?

      After you've installed the hardware for your storage system, you cable the hardware.

Cable the hardware - AFF C30 and AFF C60

After you install your AFF C30 or AFF C60 storage system hardware, cable the controllers to the network and shelves.

Before you begin

Contact your network administrator for information about connecting the storage system to your network switches.

About this task

- The cabling graphics have arrow icons showing the proper orientation (up or down) of the cable connector pull-tab when inserting a connector into a port.

  As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it over and try again.

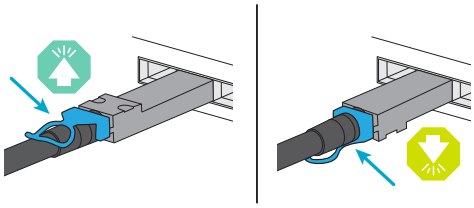

- If cabling to an optical switch, insert the optical transceiver into the controller port before cabling to the switch port.

Step 1: Cable the cluster/HA connections

Create the ONTAP cluster connections. For switchless clusters, connect the controllers to each other. For switched clusters, connect the controllers to the cluster network switches.

The cluster/HA cabling examples show common configurations.

If you do not see your configuration here, go to NetApp Hardware Universe for comprehensive configuration and slot priority information to cable your storage system.

Switchless cluster cabling

AFF C30 or AFF C60 with two 2-port 40/100 GbE I/O modules

Steps

1. Cable the Cluster/HA interconnect connections:

   The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O modules in slots 2 and 4). The ports are 40/100 GbE.

   a. Cable controller A port e2a to controller B port e2a.

   b. Cable controller A port e4a to controller B port e4a.

   I/O module ports e2b and e4b are unused and available for host network connectivity.

100 GbE Cluster/HA interconnect cables

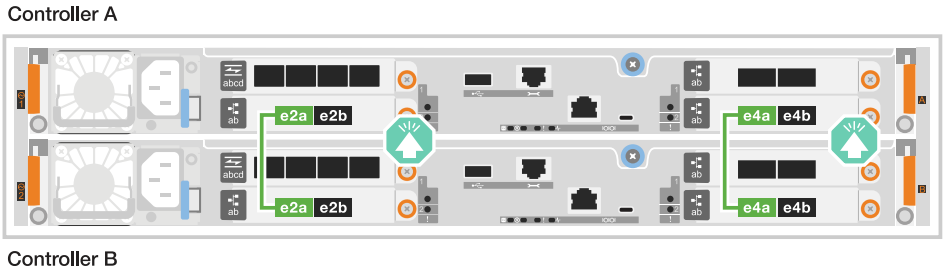

AFF C30 or AFF C60 with one 2-port 40/100 GbE I/O module

Steps

1. Cable the Cluster/HA interconnect connections:

> The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

a. Cable controller A port e4a to controller B port e4a.

b. Cable controller A port e4b to controller B port e4b.

100 GbE Cluster/HA interconnect cables

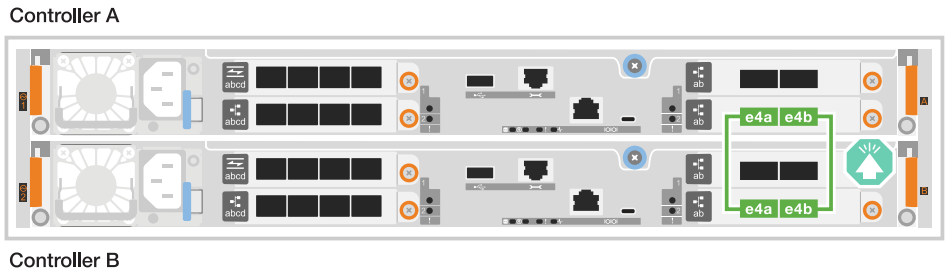

Switched cluster cabling

AFF C30 or AFF C60 with two 2-port 40/100 GbE I/O modules

Steps

1. Cable the Cluster/HA interconnect connections:

> The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O modules in slots 2 and 4). The ports are 40/100 GbE.

a. Cable controller A port e4a to cluster network switch A.

b. Cable controller A port e2a to cluster network switch B.

c. Cable controller B port e4a to cluster network switch A.

d. Cable controller B port e2a to cluster network switch B.

> I/O module ports e2b and e4b are unused and available for host network connectivity.

40/100 GbE Cluster/HA interconnect cables

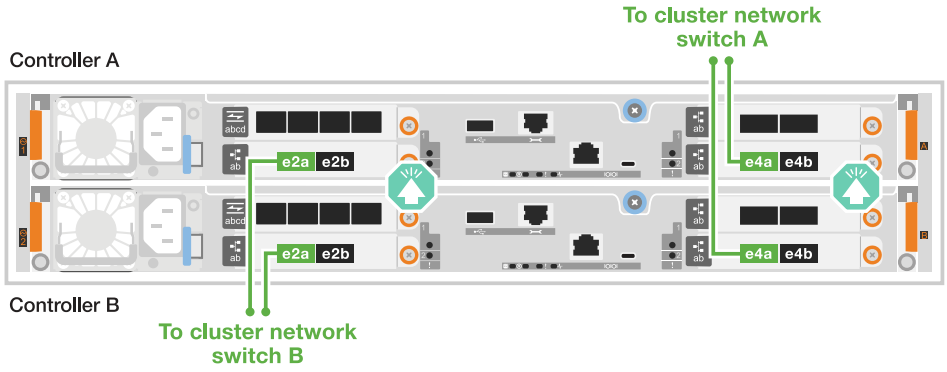

AFF C30 or AFF C60 with one 2-port 40/100 GbE I/O module

Steps

1. Cable the controllers to the cluster network switches:

> The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

a. Cable controller A port e4a to cluster network switch A.

b. Cable controller A port e4b to cluster network switch B.

c. Cable controller B port e4a to cluster network switch A.

d. Cable controller B port e4b to cluster network switch B.

40/100 GbE Cluster/HA interconnect cables

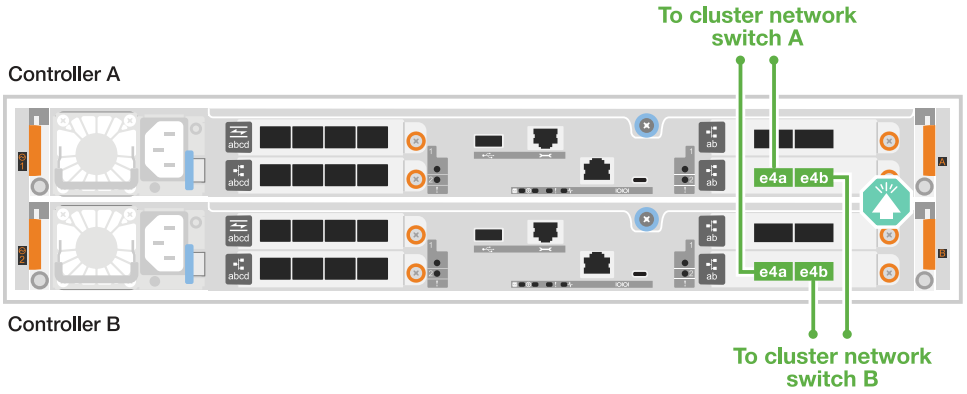

Step 2: Cable the host network connections

Cable the controllers to your Ethernet or FC host network.

> The host network cabling examples show common configurations.
>
> If you do not see your configuration here, go to NetApp Hardware Universe for comprehensive configuration and slot priority information to cable your storage system.

Ethernet host cabling

AFF C30 or AFF C60 with two 2-port 40/100 GbE I/O modules

Steps

1. On each controller, cable ports e2b and e4b to the Ethernet host network switches.

> The ports on I/O modules in slot 2 and 4 are 40/100 GbE (host connectivity is 40/100 GbE).

40/100 GbE cables

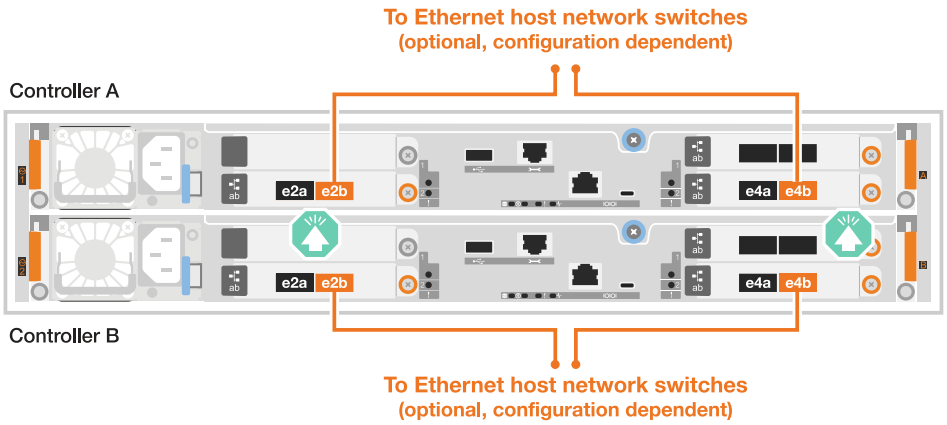

AFF C30 or AFF C60 with one 4-port 10/25 GbE I/O module

Steps

1. On each controller, cable ports e2a, e2b, e2c and e2d to the Ethernet host network switches.

10/25 GbE cables

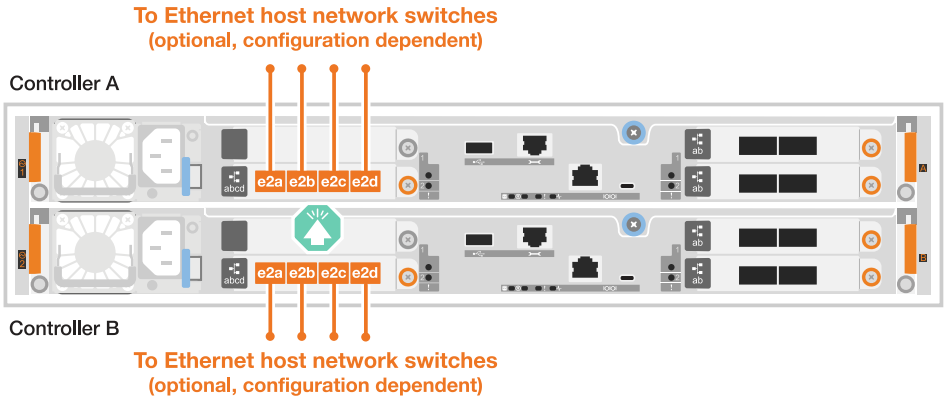

FC host cabling

AFF C30 or AFF C60 with one 4-port 64 Gb/s FC I/O module

Steps

1. On each controller, cable ports 1a, 1b, 1c and 1d to the FC host network switches.

64 Gb/s FC cables

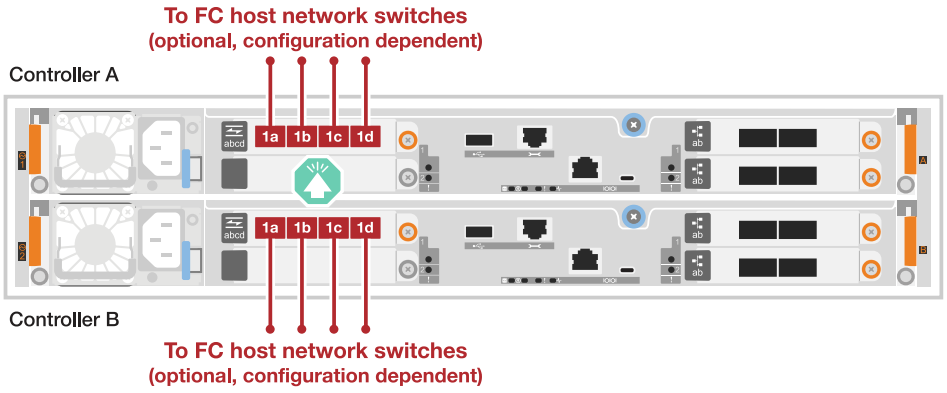

Step 3: Cable the management network connections

Cable the controllers to your management network.

1. Cable the management (wrench) ports on each controller to the management network switches.

1000BASE-T RJ-45 cables

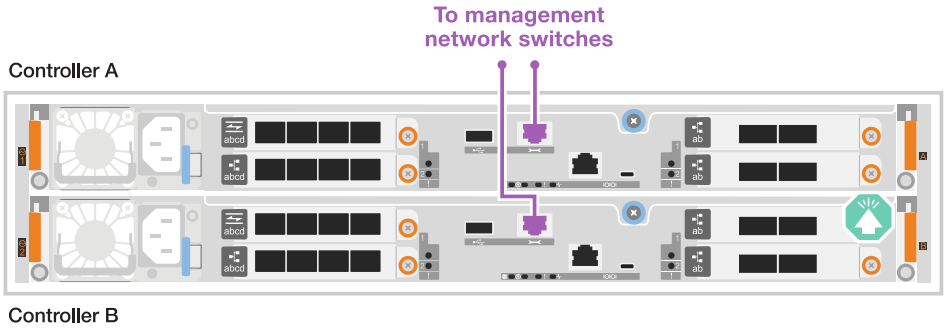

Do not plug in the power cords yet.

Step 4: Cable the shelf connections

This procedure shows you how to cable the controllers to one NS224 shelf.

About this task

- For the maximum number of shelves supported for your storage system and for all of your cabling options, such as optical and switch-attached, see NetApp Hardware Universe.
- The NS224 shelf cabling procedure shows NSM100B modules instead of NSM100 modules. The cabling is the same regardless of the type of NSM modules used, only the port names are different:
  - NSM100B modules use ports e1a and e1b on an I/O module in slot 1.
  - NSM100 modules use built-in (onboard) ports e0a and e0b.
- You cable each controller to each NSM module on the NS224 shelf using the storage cables that came with your storage system, which could be the following cable type:

100 GbE QSFP28 copper cables

- The graphics show controller A cabling in blue and controller B cabling in yellow.

Steps

1. Cable controller A to the shelf:
   a. Cable controller A port e3a to NSM A port e1a.
   b. Cable controller A port e3b to NSM B port e1b.

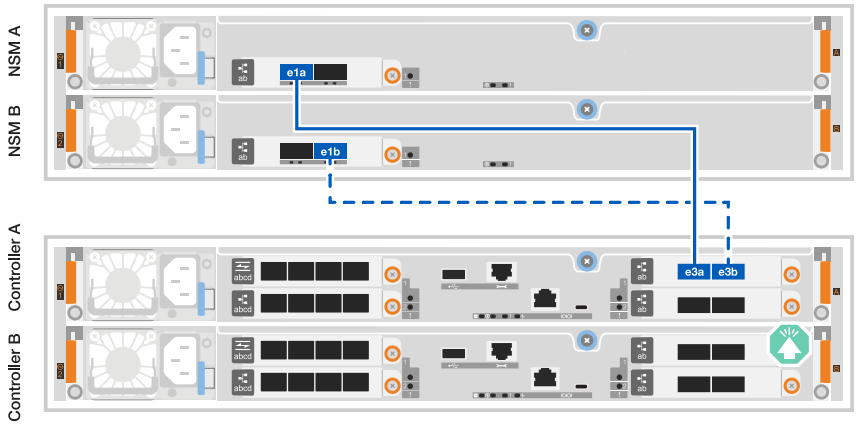

2. Cable controller B to the shelf:
   a. Cable controller B port e3a to NSM B port e1a.
   b. Cable controller B port e3b to NSM A port e1b.

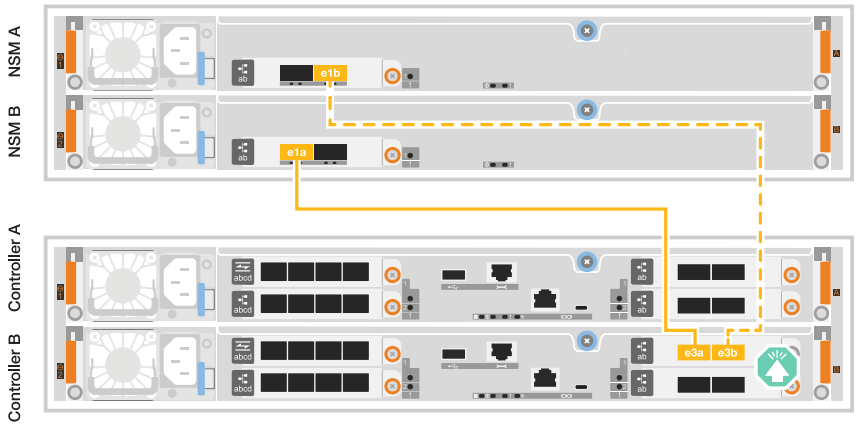

What's next?

After you've cabled the hardware for your storage system, you power on the storage system.

Power on the storage system - AFF C30 and AFF C60

After you cable the controllers to the network and shelves in your AFF C30 or AFF C60 storage system, you power on your shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

Each shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup.

Before you begin

Make sure you have a paperclip or narrow tipped ball point pen for setting NS224 storage shelf IDs.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed

shelf ID of 00.

- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, and then plug them back in) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, and then connecting the power cords to power sources on different circuits.

   The shelf powers on and boots automatically when plugged into the power source.

2. Remove the left end cap to access the shelf ID button behind the faceplate.

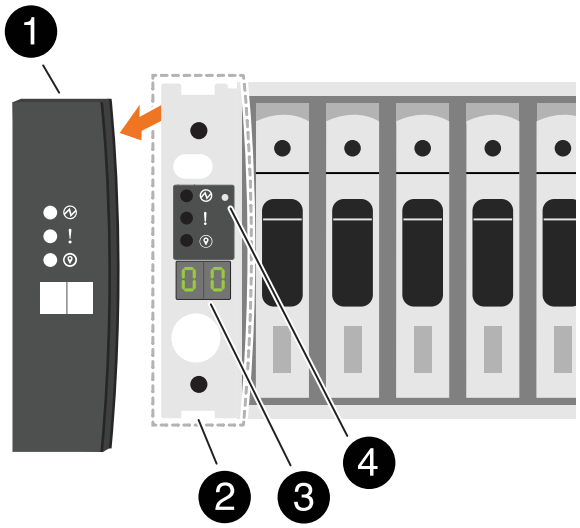

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID number
4	Shelf ID button

3. Change the first number of the shelf ID:
   a. Insert the straightened end of a paperclip or narrow tipped ball point pen into the small hole to press the shelf ID button.
   b. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

      It can take up to 15 seconds for the number to blink. This activates the shelf ID programming mode.

      If the ID takes longer than 15 seconds to blink, press and hold the shelf ID button again, making sure to press it in all the way.

   c. Press and release the shelf ID button to advance the number until you reach the desired number from

0 to 9.

Each press and release duration can be as short as one second.

The first number continues to blink.

4. Change the second number of the shelf ID:
   a. Press and hold the button until the second number on the digital display blinks.

      It can take up to three seconds for the number to blink.

      The first number on the digital display stops blinking.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

      The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the shelf ID button until the second number stops blinking.

   It can take up to three seconds for the number to stop blinking.

   Both numbers on the digital display start blinking and the amber LED illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.
   a. Unplug the power cord from both power supplies on the shelf.
   b. Wait 10 seconds.
   c. Plug the power cords back into the shelf power supplies to complete the power cycle.

      A power supply is powered on as soon as the power cord is plugged in. Its bicolored LED should illuminate green.

7. Replace the left end cap.

Step 2: Power on the controllers

After you've powered on your shelves and assigned them unique IDs, power on the storage controllers.

Steps

1. Connect your laptop to the serial console port. This will allow you to monitor the boot sequence when the controllers are powered on.
   a. Set the serial console port on the laptop to 115,200 baud with N-8-1.

      

      See your laptop's online help for instructions on how to configure the serial console port.

   b. Using the console cable provided with your storage system, connect one end of the console cable to your laptop and the other end to the serial console port on controller A.
   c. Connect the laptop to the switch on the management subnet.

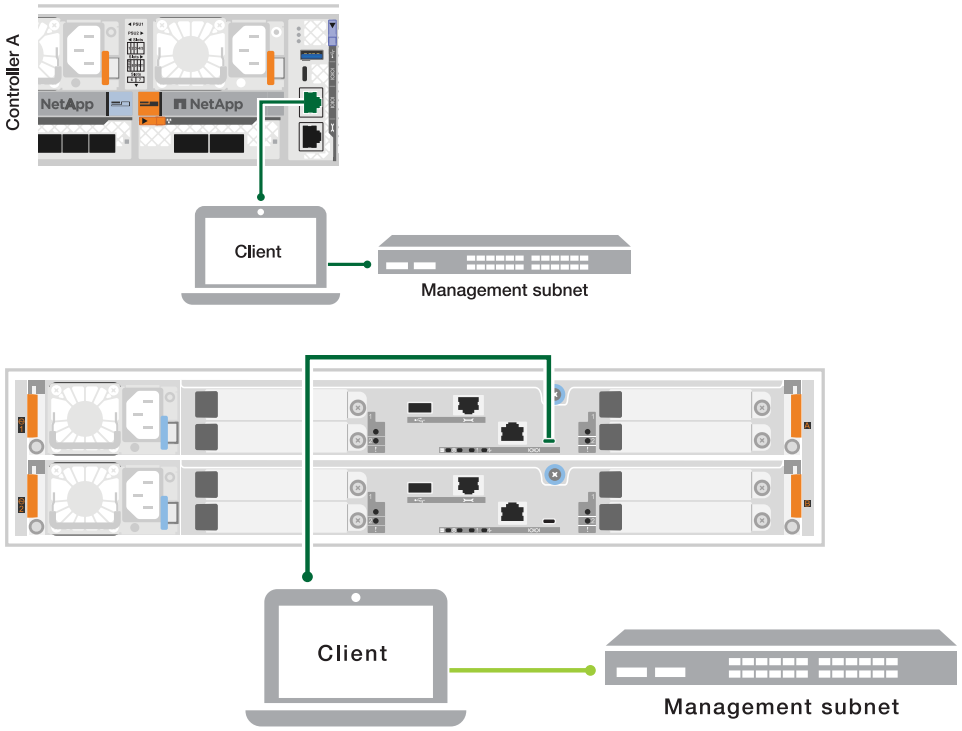

2. Assign a TCP/IP address to the laptop, using one that is on the management subnet.
3. Plug the two power cords into the controller power supplies, and then connect them to power sources on different circuits.

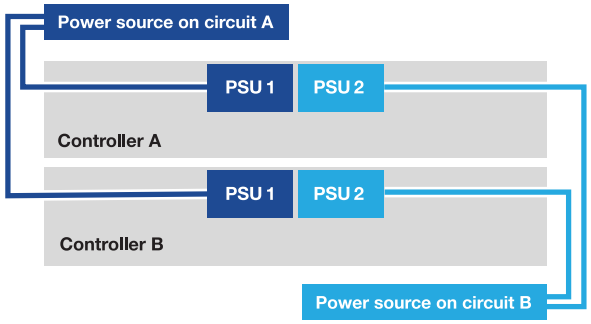

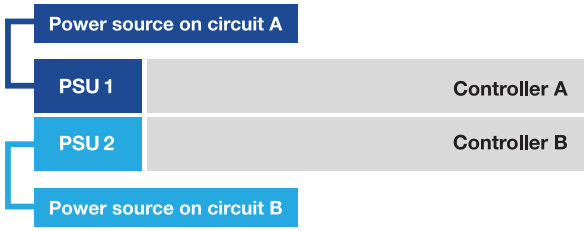

   - The system begins to boot. Initial booting might take up to eight minutes.
   - The LEDs flash on and the fans start, which indicates that the controllers are powering on.
   - The fans might be very noisy when they first start up. The fan noise during start-up is normal.
   - The shelf ID display on the front of the system chassis does not illuminate.
4. Secure the power cords using the securing device on each power supply.

What's next?

After you've powered on your storage system, you set up your cluster.

Maintain

Overview of hardware maintenance - AFF C30 and AFF C60

Maintain the hardware of your AFF C30 or AFF C60 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The procedures in this section assume that the AFF C30 or AFF C60 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF C30 and AFF C60 storage systems, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media- manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot the image from a USB drive and restore the configuration from the partner node
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller and drives.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.

NV battery	The non-volatile memory (NV) battery is responsible for providing power to the NVMEM components while data in-flight is being destaged to flash memory after a power loss.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - AFF C30 and AFF C60

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on the partner node to reinstall ONTAP on the replacement boot media in your AFF C30 or AFF C60 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

Review the boot media requirements

Review the requirements for boot media replacement.

Shut down the controller

Shut down the controller in your storage system when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the impaired controller and install the replacement boot media.

Restore the image on the boot media

Restore the ONTAP image from the partner controller.

Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automatic boot media recovery - AFF C30 and AFF C60

Before replacing the boot media in your AFF C30 or AFF C60 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0M (wrench) port on the impaired controller is working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Review the following requirements.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - AFF C30 and AFF C60

Shut down the impaired controller in your AFF C30 or AFF C60 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`)

for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - AFF C30 or AFF C60

The boot media in your AFF C30 or AFF C60 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and

then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

About this task

If needed, you can turn on the platform chassis location (blue) LEDs to aid in physically locating the affected platform. Log into the BMC using SSH and enter the `system location-led on` command.

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

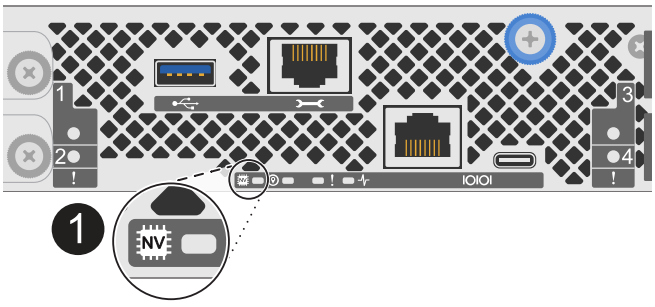

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

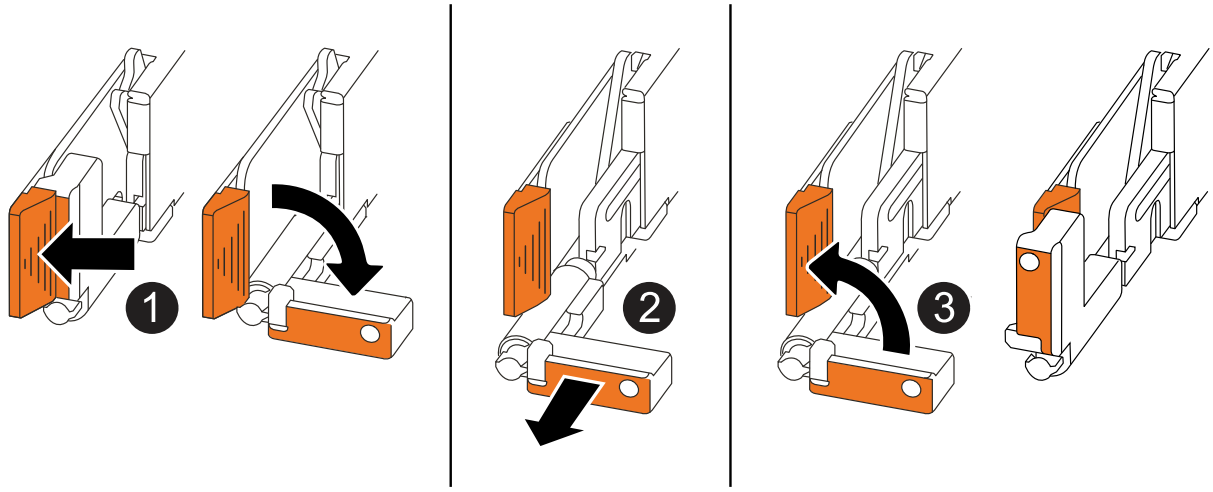

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.

7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:

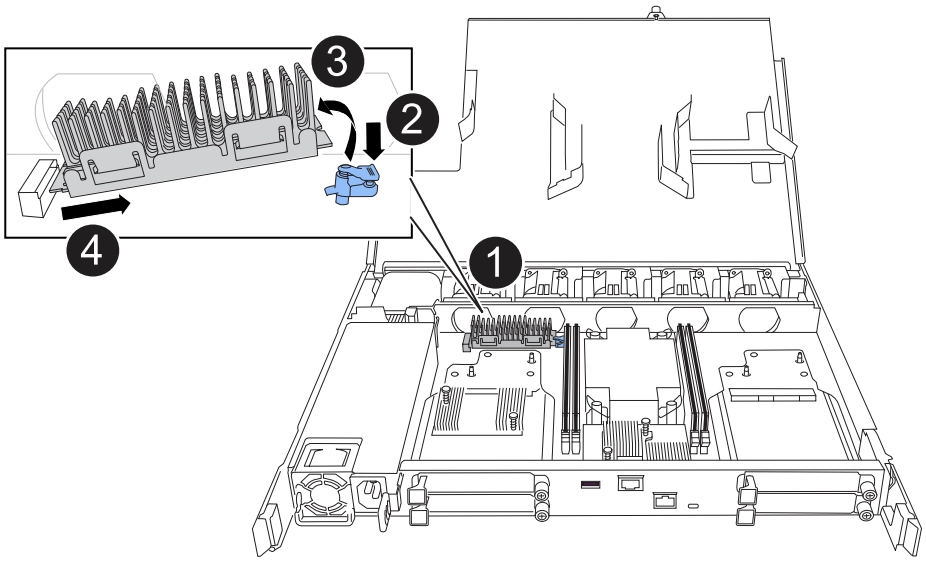

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:
   a. Remove the boot media from its package.
   b. Slide the socket end of the boot media into its socket.
   c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

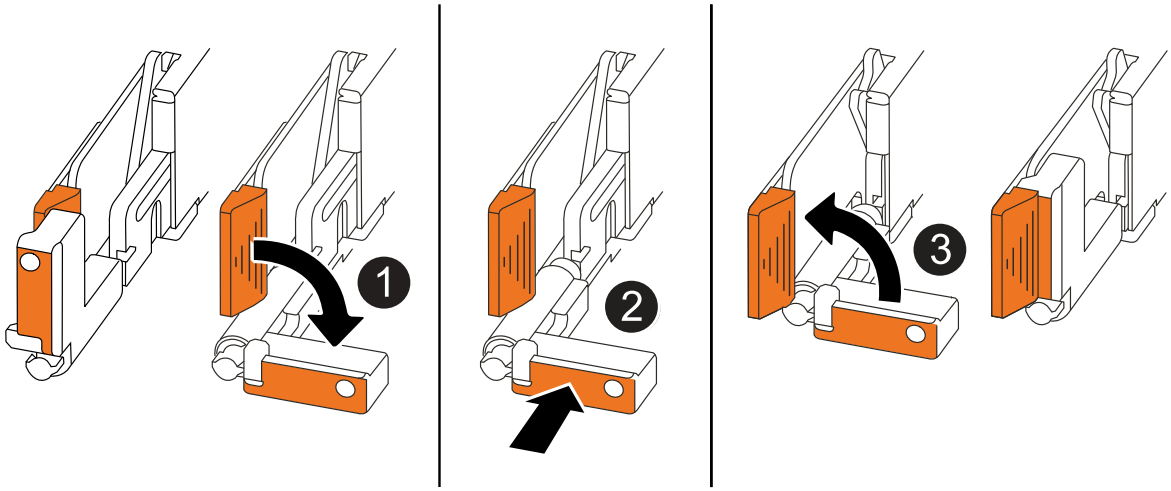

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.

   Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

      The controller boots to the LOADER prompt when fully seated in the chassis. It gets its power from the partner controller.

b. Rotate the controller handles up and lock in place with the tabs.

5. Reconnect the power cord to the PSU on the impaired controller.

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - AFF C30 and AFF C80

After installing the new boot media device in your AFF C30 and AFF C80 storage system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media part to NetApp - AFF C30 and AFF C60

If a component in your AFF C30 or AFF C60 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF C30 and AFF C60

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the AFF C30 or AFF C60 storage system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

If your system is running in ONTAP 9.17.1 and later, use the automatic boot recovery procedure.

To get started, review the recovery requirements, shut down the controller, replace the boot media, use the USB drive to restore the image, and reapply encryption settings if necessary.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check onboard encryption keys

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the impaired controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the impaired controller and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONTAP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - AFF C30 and AFF C60

Before replacing the boot media in your AFF C30 or AFF C60 storage system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption support for manual boot media recovery - AFF C30 and AFF C60

To ensure data security on your AFF C30 or AFF C60 storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

What's next?

After checking the encryption key support and status on the boot media, you need to shut down the controller.

Shut down the controller for manual boot media recovery - AFF C30 and AFF C60

Shut down the impaired controller in your AFF C30 or AFF C60 storage system to prevent data loss and maintain system stability during the manual boot media recovery process.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After shutting down the controller, you need to replace the boot media.

Replace the boot media and prepare for manual boot recovery - AFF C30 and AFF C60

The boot media in your AFF C30 or AFF C60 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

About this task

If needed, you can turn on the platform chassis location (blue) LEDs to aid in physically locating the affected platform. Log into the BMC using SSH and enter the `system location-led` *`on`* command.

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led` *`off`* command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.

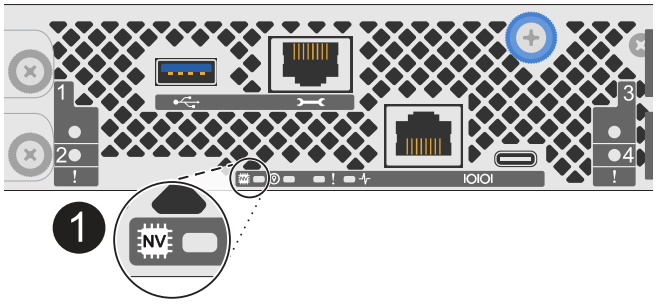

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

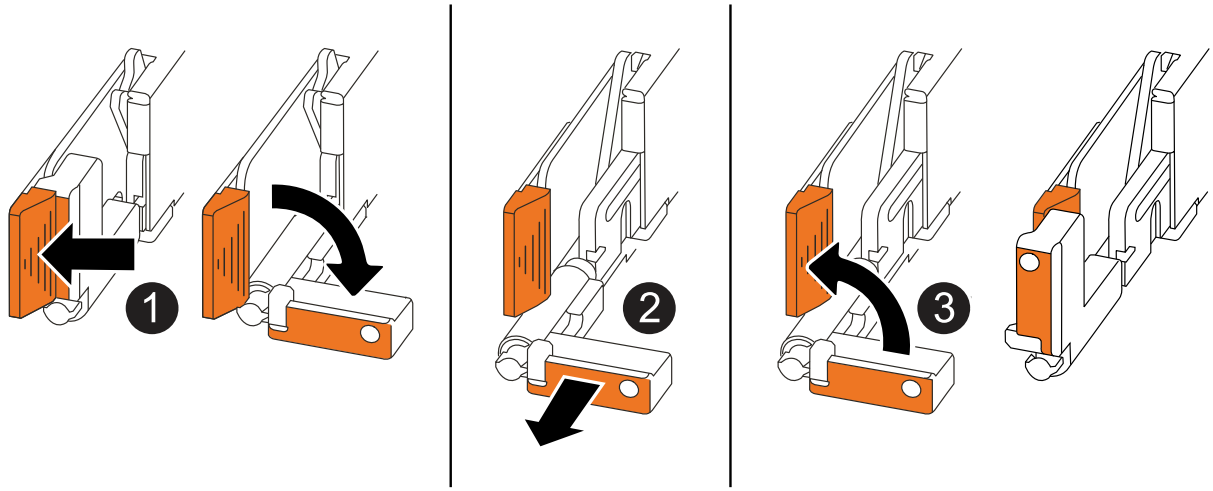

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.
7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:

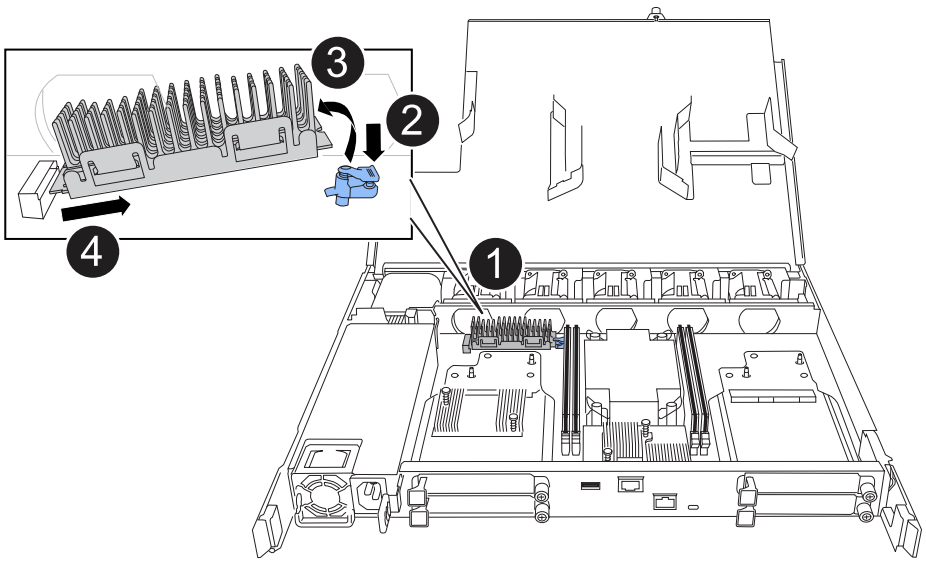

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:
   a. Remove the boot media from its package.
   b. Slide the socket end of the boot media into its socket.
   c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis, but do not reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

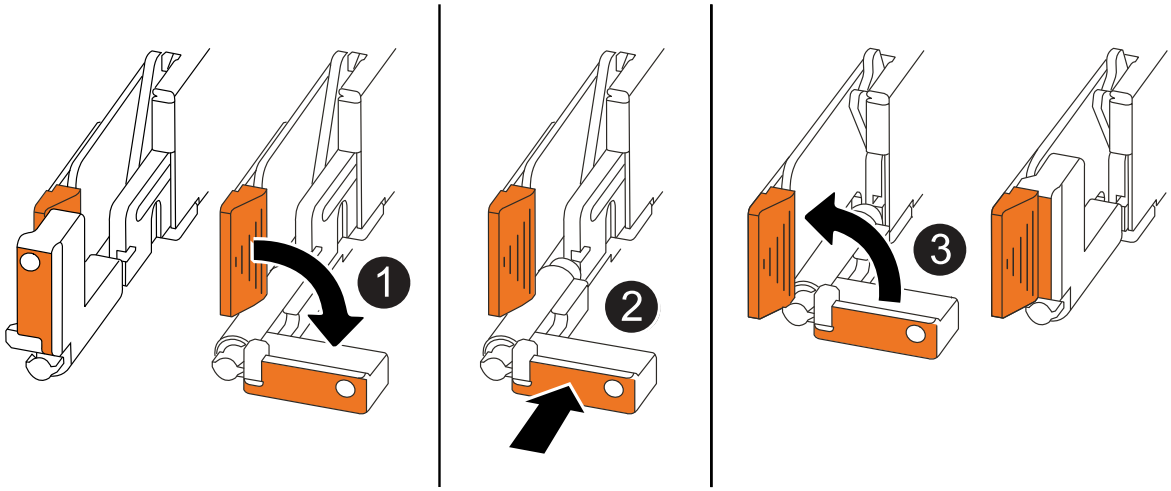

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   > Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.

   > Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

Step 4: Transfer the boot image to the boot media

The replacement boot media that you installed is without an ONTAP image so you need to transfer an ONTAP image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- You must have a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site

- If NVE is supported, download the image with NetApp Volume Encryption, as indicated in the download button.
- If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- You must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image from the Downloads link on the page, to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      The USB flash drive should have the appropriate ONTAP image of what the impaired controller is running.

   c. Remove the USB flash drive from your laptop.
2. Insert the USB flash drive into the USB-A port on the impaired controller.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

3. Fully seat the impaired controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

      The controller boots when fully seated in the chassis. It gets its power from the partner controller.

   b. Rotate the controller handles up and lock in place with the tabs.
4. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

5. Reconnect the power cord to the power supply (PSU) on the impaired controller.

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.

If you are reconnecting a…	Then…
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

After replacing the boot media, you need to boot the recovery image.

Manual boot media recovery from a USB drive - AFF C30 and AFF C60

After installing the new boot media device in your AFF C30 or AFF C60 storage system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `...would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

What's next?

After booting the recovery image, you need to restore encryption on the boot media.

Restore encryption keys after manual boot recovery - AFF C30 and AFF C60

Restore encryption on the replacement boot media in your AFF C30 or AFF C60 storage system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    `security key-manager key query -restored false.`

    ```
    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    `storage failover giveback -fromnode local`

13. Restore automatic giveback, if you disabled it, by entering the following command.

    `storage failover modify -node local -auto-giveback true`

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    `system node autosupport invoke -node * -type all -message MAINT=END`

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
**********************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
**********************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

What's next?

After restoring encryption on the boot media, you need to return the failed part to NetApp.

Return the failed part to NetApp - AFF C30 and AFF C60

If a component in your AFF C30 or AFF C60 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - AFF C30 and AFF C60

Get started with replacing the chassis of your AFF C30 or AFF C60 storage system by reviewing the replacement requirements, shutting down the controllers, replacing the chassis, and verifying system operations.

1 Review the chassis replace requirements

To replace the chassis, you must meet certain requirements.

2 Shut down the controllers

Shut down the controllers so you can perform maintenance on the chassis.

3 Replace the chassis

Replacing the chassis includes moving the drives and any drive blanks, controllers (with the power supplies), and bezel from the impaired chassis to the new chassis, and swapping out the impaired chassis with the new chassis of the same model as the impaired chassis.

4 Complete chassis replacement

Verify the HA state of the chassis and return the failed part to NetApp.

Requirements to replace the chassis - AFF C30 and AFF C60

Before replacing the chassis of your AFF C30 or AFF C60 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement chassis, and the necessary tools.

Review the following requirements and considerations.

Requirements

- The replacement chassis must be the same model as the impaired chassis. This procedure is for a like-for-like replacement, not for an upgrade.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Considerations

- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your storage system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, drives, any drive blanks, and controllers to the new chassis.

What's next?

After you've reviewed the requirements to replace the chassis, you need to shut down the controllers

Shut down the controllers - AFF C30 and AFF C60

Shut down the controllers in your AFF C30 or AFF C60 storage system to prevent data loss and ensure system stability when replacing the chassis.

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true
   ```

For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

What's next?

After you've shut down the controllers, you need to replace the chassis.

Replace the chassis - AFF C30 and AFF C60

Replace the chassis of your AFF C30 or AFF C60 storage system when a hardware failure requires it. The replacement process involves removing the controllers, removing the drives, installing the replacement chassis, and reinstalling the chassis components.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

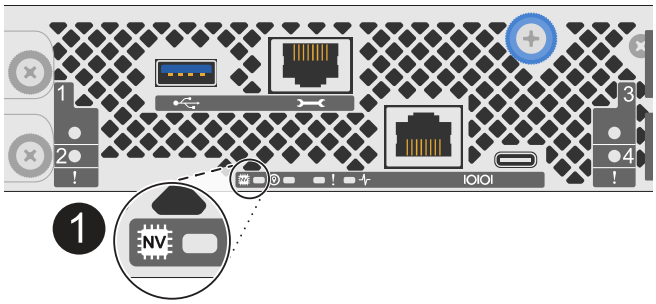

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

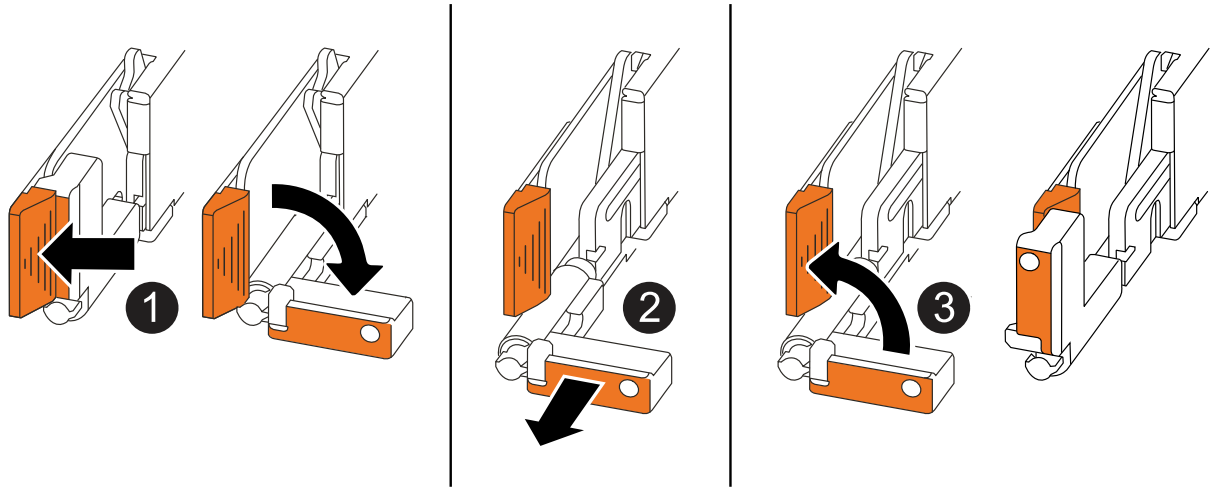

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Repeat these steps for the other controller in the chassis.

Step 2: Remove the drives from the impaired chassis

You need to remove all of the drives and any drive blanks from the impaired chassis so that later in the procedure you can install them in the replacement chassis.

1. Gently remove the bezel from the front of the storage system.
2. Remove the drives and any drive blanks:

Keep track of what drive bay each drive and drive blank was removed from because they must be installed in the same drive bays in the replacement chassis.

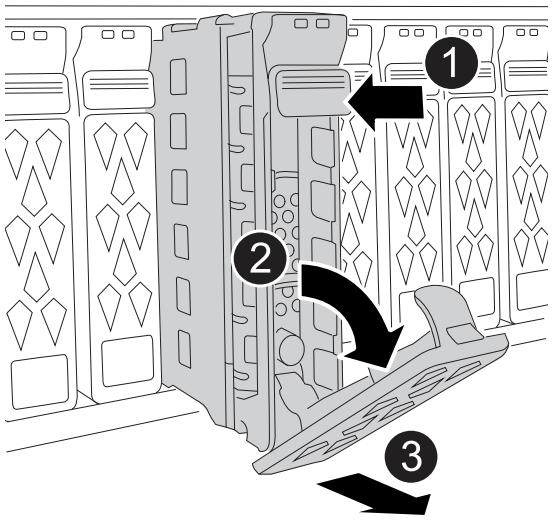

1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.

3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

3. Set the drives aside on a static-free cart or table.

Step 2: Replace the chassis from within the equipment rack or system cabinet

You remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis, install the drives, any drive blanks, and then install the bezel.

1. Remove the screws from the impaired chassis mount points.

   Set the screws aside to use later in this procedure.

   If the storage system shipped in a NetApp system cabinet, you must remove additional screws at the rear of the chassis before the chassis can be removed.

2. Using two people or a power lift, remove the impaired chassis from the equipment rack or system cabinet by sliding it off the rails, and then set it aside.
3. Using two people, install the replacement chassis into the equipment rack or system cabinet by sliding it onto the rails.
4. Secure the front of the replacement chassis to the equipment rack or system cabinet using the screws you removed from the impaired chassis.

Step 4: Install the controllers and drives

Install the controllers and drives into the replacement chassis and reboot the controllers.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when installing a controller, and can be used as a reference for the rest of the controller installation steps.

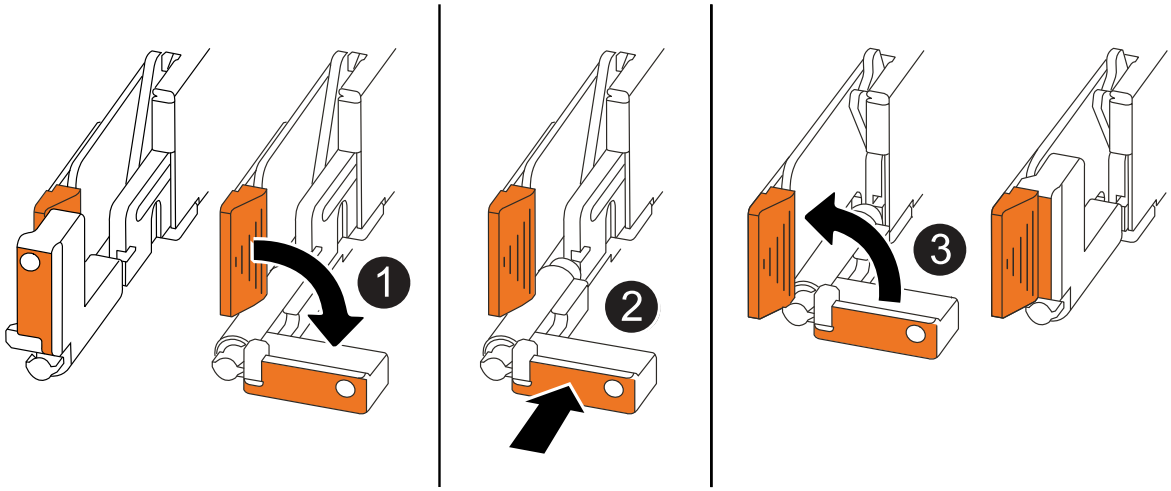

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis and push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

1. Insert one of the controllers into the chassis:
   a. Align the back of the controller with the opening in the chassis.
   b. Firmly push on the handles until the controller meets the midplane and is fully seated in the chassis.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   c. Rotate the controller handles up and lock in place with the tabs.
2. Recable the controller, as needed, except for the power cords.
3. Repeat these steps to install the second controller into the chassis.
4. Install the drives and any drive blanks you removed from the impaired chassis into the replacement chassis:

   The drives and drive blanks must be installed in the same drive bays in the replacement chassis.

   a. With the cam handle in the open position, use both hands to insert the drive.
   b. Gently push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

   d. Repeat the process for the remaining drives.

5. Install the bezel.
6. Reconnect the power cords to the power supplies (PSU) in the controllers.

   Once power is restored to a PSU, the status LED should be green.

   The controllers begin to boot as soon as the power is restored.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. If controllers boot to the LOADER prompt, reboot the controllers:

   ```
   boot_ontap
   ```

8. Turn AutoSupport back on:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

What's next?

After you've replaced the impaired chassis and reinstalled the components into it, you need to complete the chassis replacement.

Complete chassis replacement - AFF C30 and AFF C60

Verify the HA state of the chassis and then return the failed part to NetApp to complete the final step in the AFF C30 and AFF C60 chassis replacement procedure.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your storage system configuration.

1. In Maintenance mode, from either controller, display the HA state of the local controller and chassis:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your storage system configuration:
   a. Set the HA state for the chassis:

      ```
      ha-config modify chassis HA-state
      ```

      The value for HA-state should be *ha*.

The value for HA-state can be one of the following:
* ***ha***
* *mcc* (not supported in ASA)

b. Confirm that the setting has changed:

```
ha-config show
```

3. If you have not already done so, recable the rest of your storage system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - AFF C30 and AFF C60

Get started with replacing the controller in your AFF C30 or AFF C60 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

Review the controller replacement requirements

To replace the controller, you must meet certain requirements.

Shut down the impaired controller

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

Replace the controller

Replacing the controller includes removing the impaired controller, moving FRU components to the replacement controller, installing the replacement controller in the chassis, setting the time and date, and then recabling.

Restore and verify the system configuration

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Give back the controller

Transfer the ownership of storage resources back to the replacement controller.

6 Complete controller replacement

Verify the LIFs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - AFF C30 and AFF C60

Before replacing the controller in your AFF C30 or AFF C60 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements and considerations for the controller replacement procedure.

Requirements

- All shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.
- You must replace a controller with a controller of the same model type. You cannot upgrade your system by just replacing the controller.
- You cannot change any drives or shelves as part of this procedure.
- You must always capture the controller's console output to a text log file.

  The console output provides you with a record of the procedure you can use to troubleshoot issues you might encounter during the replacement process.

Considerations

It is important that you apply the commands in this procedure to the correct controller:

- The *impaired* controller is the controller that is being replaced.
- The *replacement* controller is the new controller that is replacing the impaired controller.
- The *healthy* controller is the surviving controller.

What's next?

After you've reviewed the requirements to replace the impaired controller, you need to shut down the impaired controller.

Shut down the impaired controller - AFF C30 and AFF C60

Shut down the impaired controller in your AFF C30 or AFF C60 storage system to prevent data loss and ensure system stability when replacing the controller.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After you've shut down the impaired controller, you need to replace the controller.

Replace the controller - AFF C30 and AFF C60

Replace the controller in your AFF C30 or AFF C60 storage system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.

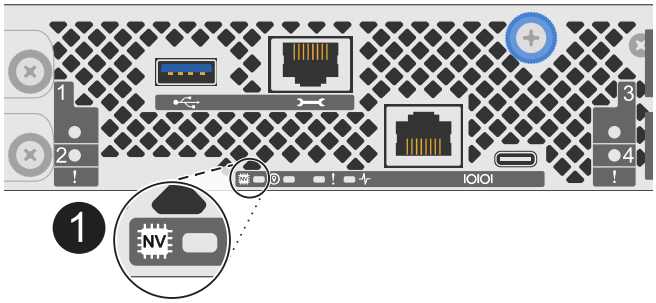

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

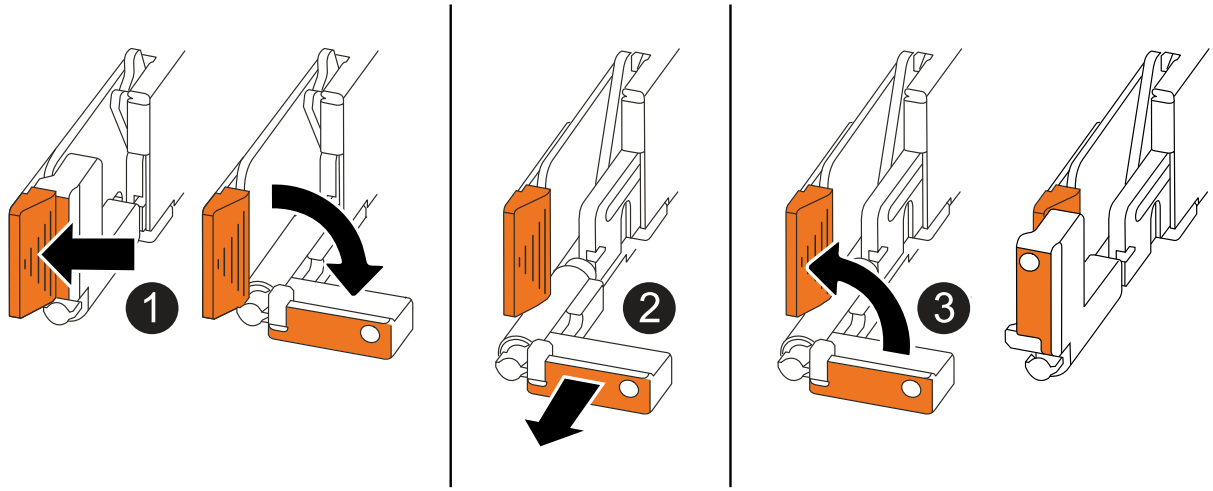

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Move the power supply

Move the power supply (PSU) to the replacement controller.

1. Move the PSU from the impaired controller:

   Make sure the left side controller handle is in the upright position to allow you access to the PSU.

Option 1: Move an AC PSU

To move an AC PSU, complete the following steps.

Steps

1. Remove the AC PSU from the impaired controller:

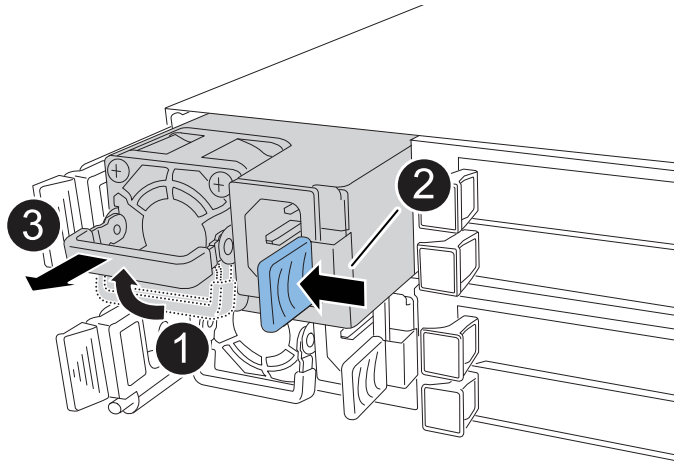

1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

2. Insert the PSU into the replacement controller:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.

Option 2: Move a DC PSU

To move a DC PSU, complete the following steps.

Steps

1. Remove the DC PSU from the impaired controller:

a. Rotate the handle up, to its horizontal position, and then grasp it.

b. With your thumb, press the terracotta tab to release the locking mechanism.

c. Pull the PSU out of the controller while using your other hand to support its weight.

> The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.

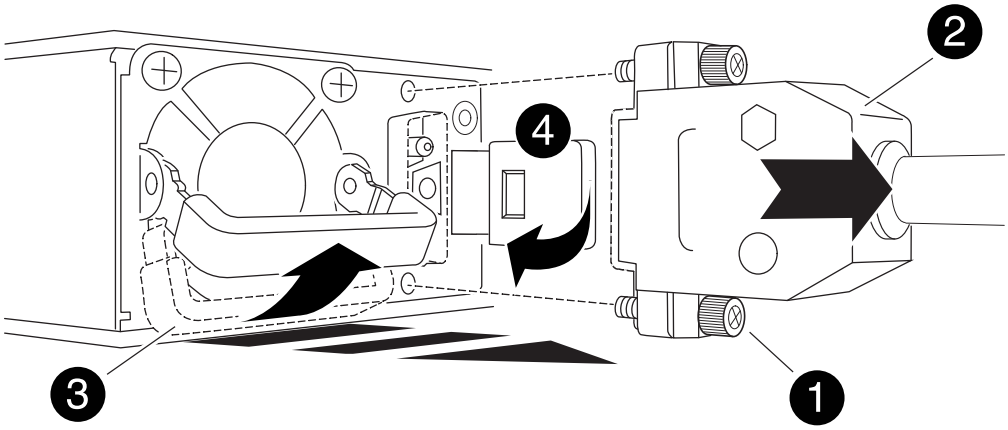

1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

2. Insert the PSU into the replacement controller:

a. Using both hands, support and align the edges of the PSU with the opening in the controller.

b. Gently slide the PSU into the controller until the locking tab clicks into place.

A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

> To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

c. Rotate the handle down, so it is out of the way of normal operations.

Step 3: Move the fans

Move the fans to the replacement controller.

1. Remove one of the fans from the impaired controller:

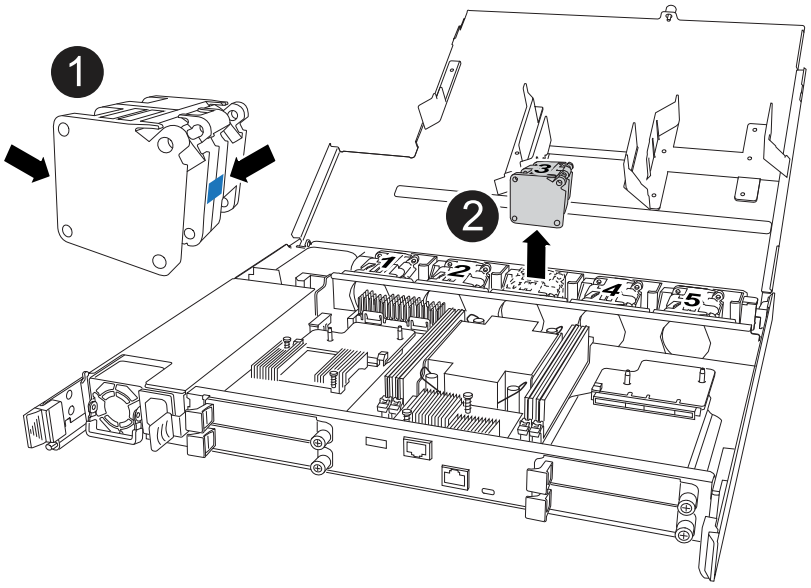

1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

2. Insert the fan into the replacement controller by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.
3. Repeat these steps for the remaining fans.

Step 4: Move the NV battery

Move the NV battery to the replacement controller.

1. Remove the NV battery from the impaired controller:

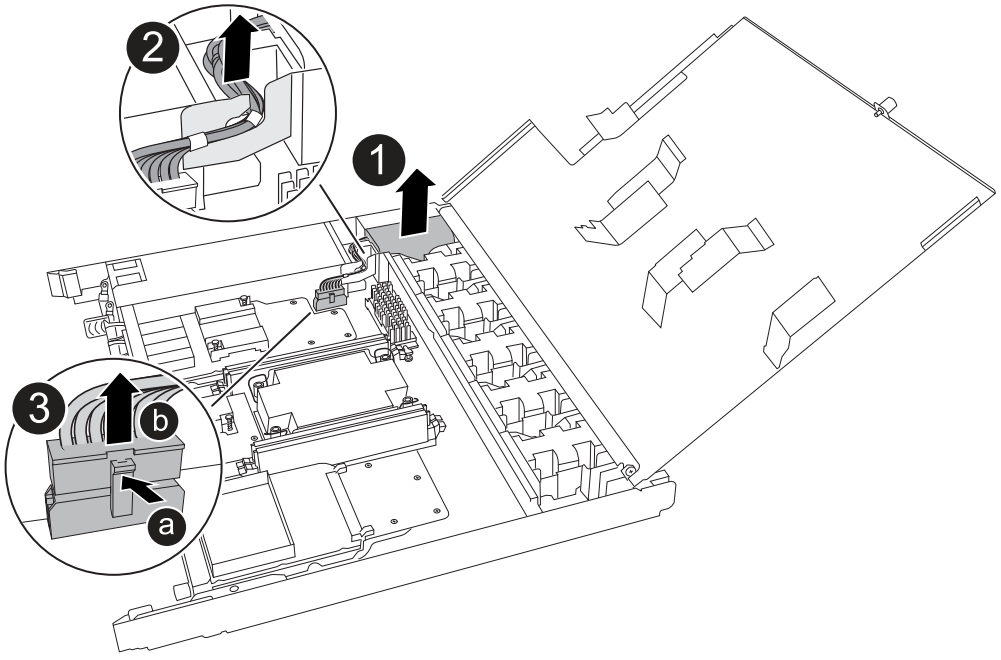

1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

2. Install the NV battery into the replacement controller:
   a. Plug the wiring connector into its socket.
   b. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   c. Place the NV battery into the compartment.

      The NV battery should sit flush in its compartment.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller.

If you have DIMM blanks, you do not need to move them, the replacement controller should come with them installed.

1. Remove one of the DIMMs from the impaired controller:

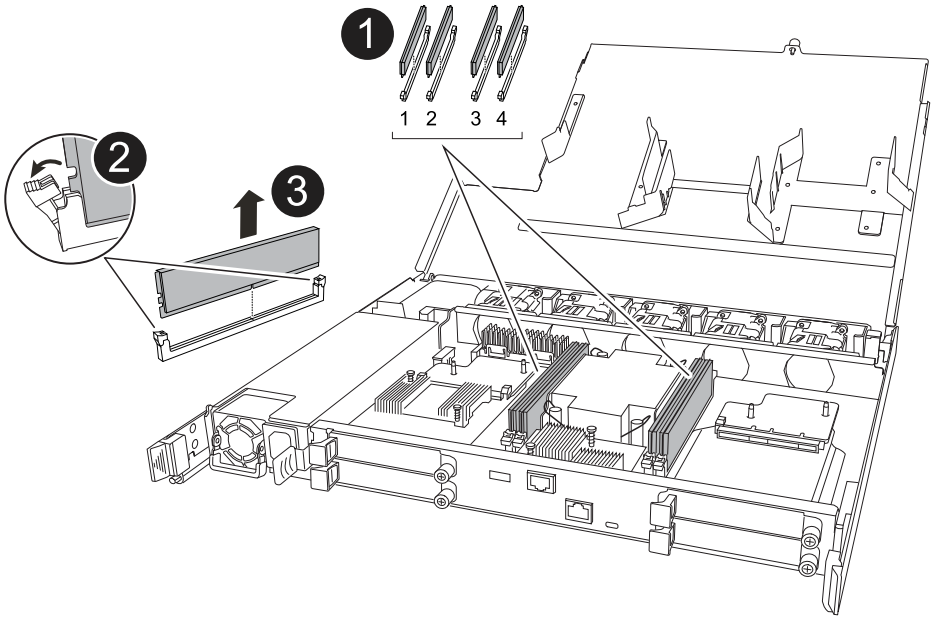

1	DIMM slot numbering and positions. Depending on your storage system model, you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller in the proper orientation. • Eject the DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

2. Install the DIMM in the replacement controller:
   a. Make sure that the DIMM ejector tabs on the connector are in the open position.
   b. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

      When inserted correctly, the DIMM goes in easily but fits tightly in the slot. If not, reinsert the DIMM.

   c. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
   d. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.
3. Repeat these steps for the remaining DIMMs.

Step 6: Move the boot media

Move the boot media to the replacement controller.

1. Remove the boot media from the impaired controller:

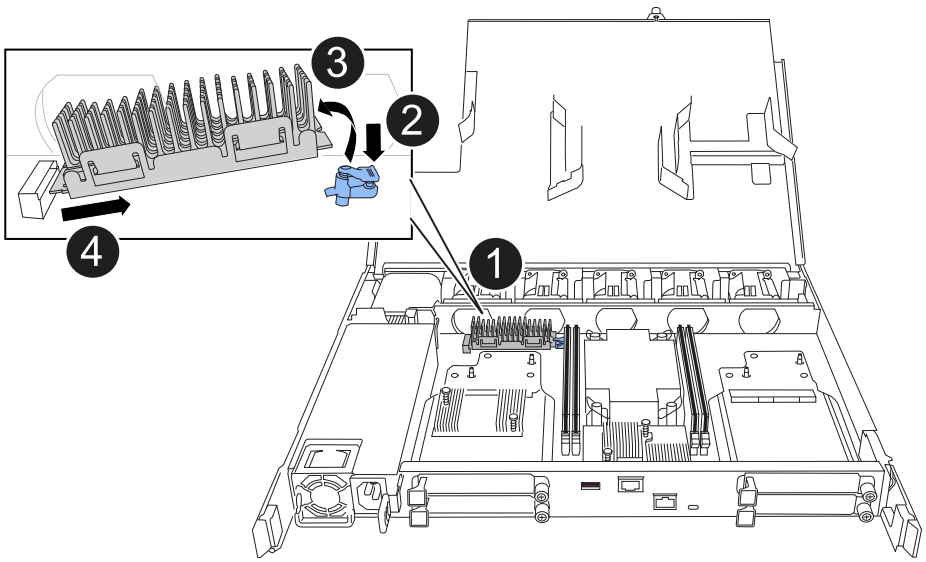

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

2. Install the boot media into the replacement controller:
   a. Slide the socket end of the boot media into its socket.
   b. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 7: Move the I/O modules

Move the I/O modules and any I/O blanking modules to the replacement controller.

1. Unplug cabling from one of the I/O modules.

   Make sure to label the cables so that you know where they came from.

2. Remove the I/O module from the impaired controller:

   Make sure that you keep track of which slot the I/O module was in.

   If you are removing the I/O module in slot 4, make sure the right side controller handle is in the upright position to allow you access to the I/O module.

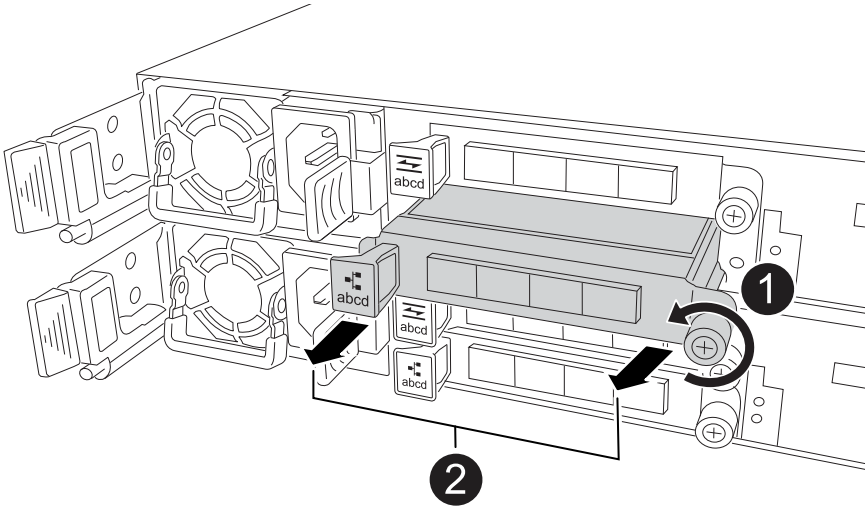

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

3. Install the I/O module into the replacement controller:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
4. Repeat these steps to move the remaining I/O modules and any I/O blanking modules to the replacement controller.

Step 8: Install the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

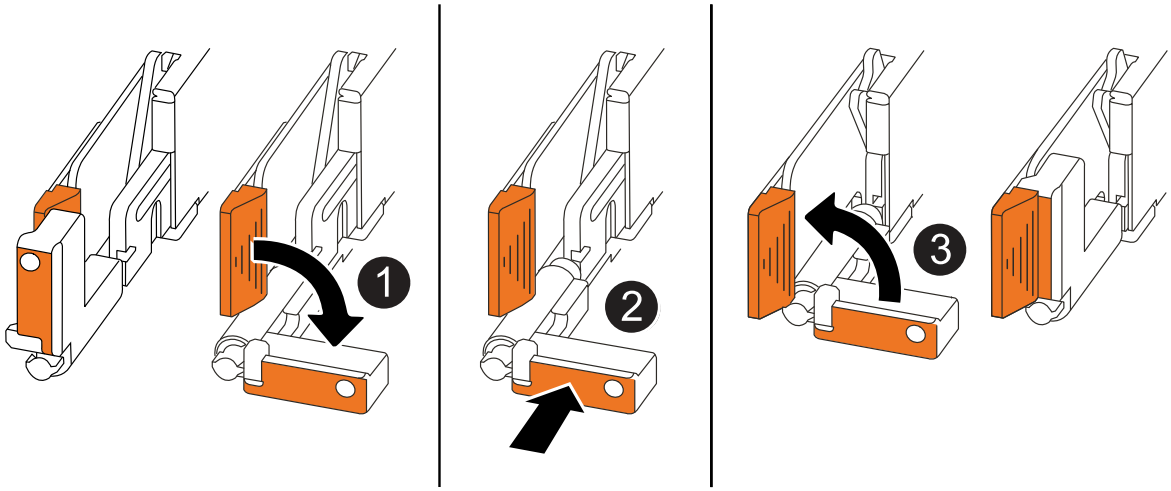

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Take the controller to the LOADER prompt by pressing CTRL-C to abort AUTOBOOT.
6. Set the time and date on the controller:

   Make sure you are at the controller's LOADER prompt.

   a. Display the date and time on the controller:

      ```
      show date
      ```

      Time and date default is in GMT. You have the option to display in local time and in 24hr mode.

   b. Set the current time in GMT:

      ```
      set time hh:mm:ss
      ```

      You can get the current GMT from the healthy node:

      ```
      date -u
      ```

   c. Set the current date in GMT:

      ```
      set date mm/dd/yyyy
      ```

      You can get the current GMT from the healthy node:
      ```
      date -u
      ```
7. Recable the controller as needed.
8. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

After you've replaced the impaired controller, you need to restore the system configuration.

Restore and verify the system configuration - AFF C30 and AFF C60

Verify that the controller's HA configuration is active and functioning correctly in your AFF C30 or AFF C60 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

You must verify the `HA` state of the controller and, if necessary, update the state to match your storage system configuration.

1. Boot to maintenance mode:

   ```
   boot_ontap maint
   ```

   a. Enter `y` when you see *Continue with boot?*.

      If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

5. If the displayed system state of the controller does not match your storage system configuration, set the `HA` state for the controller:

   ```
   ha-config modify controller ha
   ```

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed:

   ```
   ha-config show
   ```

Step 2: Verify disk list

1. Verify that the adapter lists the paths to all disks:

   ```
   storage show disk -p
   ```

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode:

```
halt
```

What's next?

After you've restored and verified your system configuration, you need to give back the controller.

Give back the controller - AFF C30 and AFF C60

Return control of storage resources to the replacement controller so your AFF C30 or AFF C60 storage system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption, Onboard Key Manager (OKM) encryption, or External Key Manager (EKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`

   If you encounter errors, contact NetApp Support.

9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

c. Verify status of the keys: `security key-manager key query -restored false`

   The output should show no results when when properly synchronized.

   If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

External key manager (EKM)

Reset encryption and return the controller to normal operation.

Steps

1. If the root volume is encrypted with External Key Manager and the console cable is connected to the replacement node, enter `boot_ontap menu` and select option `11`.
2. If these questions appear, answer `y` or `n` as appropriate:

   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}

   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}

   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}

   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}

   Do you know the KMIP server address? {y/n}

   Do you know the KMIP port? {y/n}

   Contact NetApp Support if you have issues.

3. Supply the information for:
   - The client certificate (client.crt) file contents
   - The client key (client.key) file contents
   - The KMIP server CA(s) (CA.pem) file contents
   - The IP address for the KMIP server
   - The port for the KMIP server
4. Once the system processes, you see the Boot Menu. Select '1' for normal boot.
5. Check the takeover status: `storage failover show`
6. Ensure any core dumps on the repaired node are saved by going to advanced mode `set -privilege advanced` and then `run local partner nosavecore`.
7. Return the impaired controller to normal operation by giving back its storage: `storage failover`

`giveback -ofnode impaired_node_name`

8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
9. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

After you've transferred the ownership of storage resources to the replacement controller, you need to complete the controller replacement procedure.

Complete controller replacement - AFF C30 and AFF C60

To complete the controller replacement for your AFF C30 or AFF C60 storage system, first restore the NetApp Storage Encryption configuration (if necessary) and install the required licenses on the new controller. Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, register the new controller's serial number and then return the failed part to NetApp.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

  Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The licenses keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.
- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs, register the serial number, and check cluster health

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF C30 and AFF C60

Replace a DIMM in your AFF C30 or AFF C60 storage system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- All other components in the storage system must be working correctly; if not, contact NetApp Support before continuing.

- You must replace the failed FRU component with a replacement FRU component you received from your provider.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

   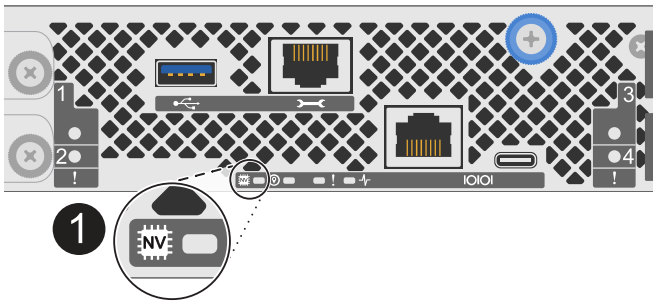

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

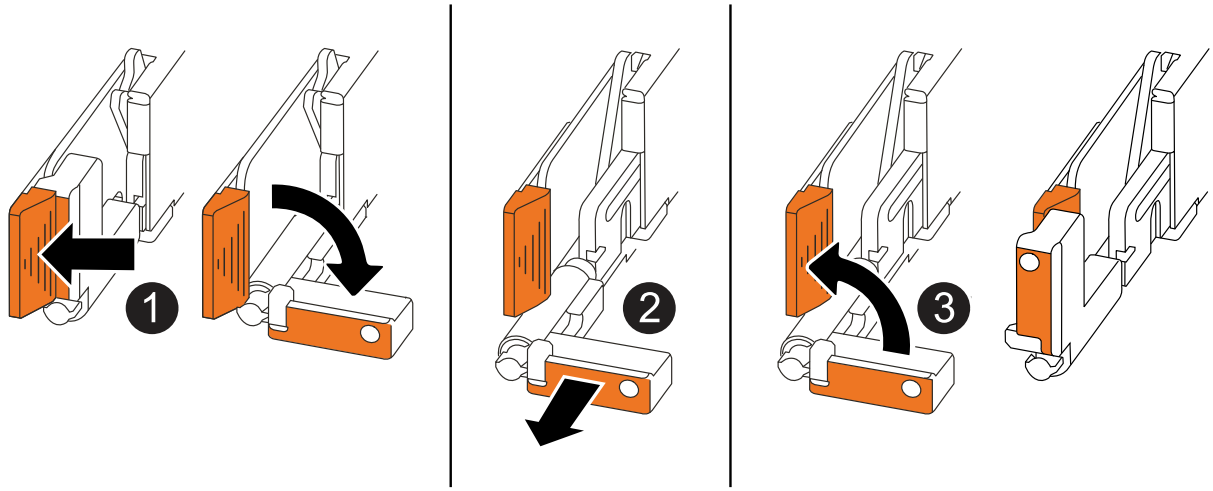

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace a DIMM

To replace a DIMM, locate the faulty DIMM inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller and identify the faulty DIMM.

> Consult either the Netapp Hardware Universe or the FRU map on the cover of the controller for exact DIMM locations.

3. Remove the faulty DIMM:

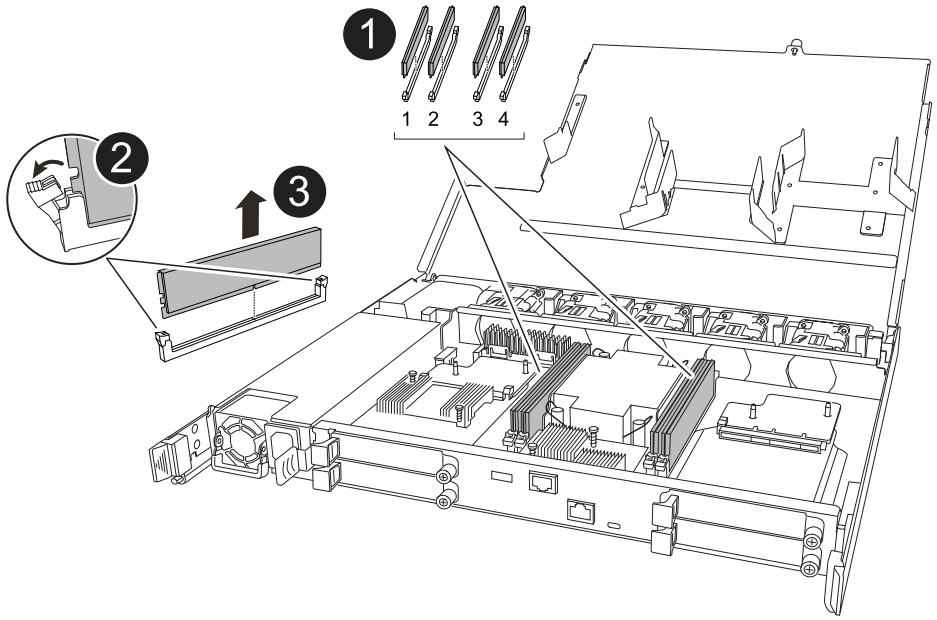

1	DIMM slot numbering and positions. Depending on your storage system model you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM using the same orientation. • Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

4. Install the replacement DIMM:
   a. Remove the replacement DIMM from its antistatic shipping bag.
   b. Make sure that the DIMM ejector tabs on the connector are in the open position.
   c. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

      When inserted correctly, the DIMM goes in easily but fits tightly in the slot. Reinsert the DIMM if you feel it is not inserted correctly.

   d. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
   e. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

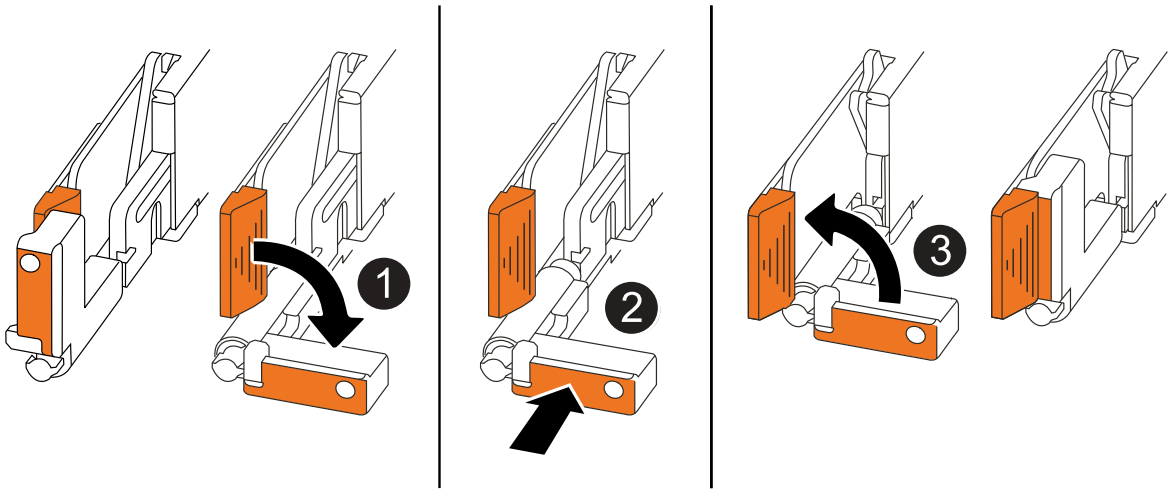

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a drive - AFF C30 and AFF C60

Replace a drive in your AFF C30 or AFF C60 storage system when a drive fails or requires an upgrade. The replacement process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed SSD drive nondisruptively while I/O is in progress.

Before you begin

- The drive that you are installing must be supported by your storage system.

  NetApp Hardware Universe

- If self-encrypting drive (SED) authentication is enabled, you must use the SED replacement instructions in the ONTAP documentation.

  Instructions in the ONTAP documentation describe additional steps you must perform before and after replacing an SED.

  NetApp encryption overview with the CLI

- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.
- Verify that the drive you are removing is failed.

  You can verify that the drive is failed by running the `storage disk show -broken` command. The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  Depending on the drive type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

About this task

- When replacing a failed drive, you must wait 70 seconds between the removal of the drive and the insertion of the replacement drive to allow the storage system to recognize that a drive was removed.
- The best practice is to have the current version of the Disk Qualification Package (DQP) installed before hot-swapping a drive.

  Having the current version of the DQP installed allows your system to recognize and use newly qualified drives. This avoids system event messages about having noncurrent drive information and prevention of drive partitioning because drives are not recognized. The DQP also notifies you of noncurrent drive firmware.

  NetApp Downloads: Disk Qualification Package

- The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on

your system before replacing FRU components.

NetApp Downloads: Disk Shelf Firmware

NetApp Downloads: Disk Drive Firmware

> Do not revert firmware to a version that does not support your shelf and its components.

- Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

  > Drive firmware checks occur every two minutes.

- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment if it is enabled.

   > You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled:

      ```
      storage disk option show
      ```

      You can enter the command on either controller.

      If automatic drive assignment is enabled, the output shows `on` in the `Auto Assign` column (for each controller).

   b. If automatic drive assignment is enabled, disable it:

      ```
      storage disk option modify -node node_name -autoassign off
      ```

      You must disable automatic drive assignment on both controllers.

2. Properly ground yourself.
3. Remove the bezel from the front of the storage system.
4. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

5. Remove the failed drive:

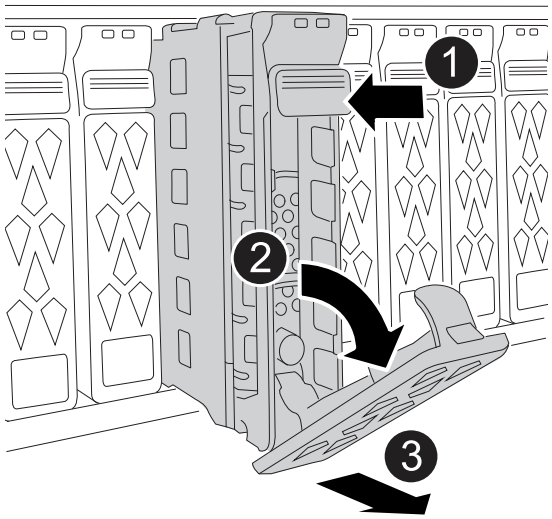

1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

6. Wait a minimum of 70 seconds before inserting the replacement drive.
7. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the drive.
   b. Gently push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
8. Verify that the drive's activity (green) LED is illuminated.

   When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

9. If you are replacing another drive, repeat steps 4 through step 8.
10. Reinstall the bezel on the front of the storage system.
11. If you disabled automatic drive assignment in step 1, manually assign drive ownership, and then reenable automatic drive assignment if needed:
    a. Display all unowned drives:

       `storage disk show -container-type unassigned`

       You can enter the command on either controller.

    b. Assign each drive:

       `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed:

       `storage disk option modify -node node_name -autoassign on`

       You must reenable automatic drive assignment on both controllers.
12. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan module - AFF C30 and AFF C60

Replace a fan module in your AFF C30 or AFF C60 storage system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

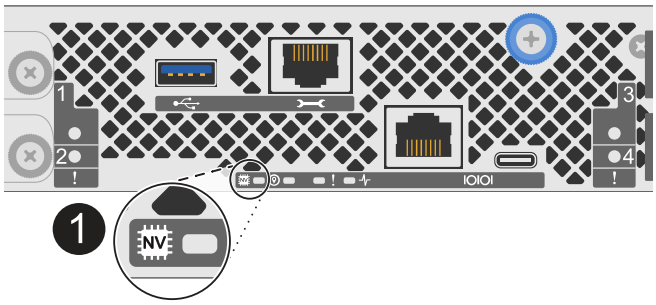

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

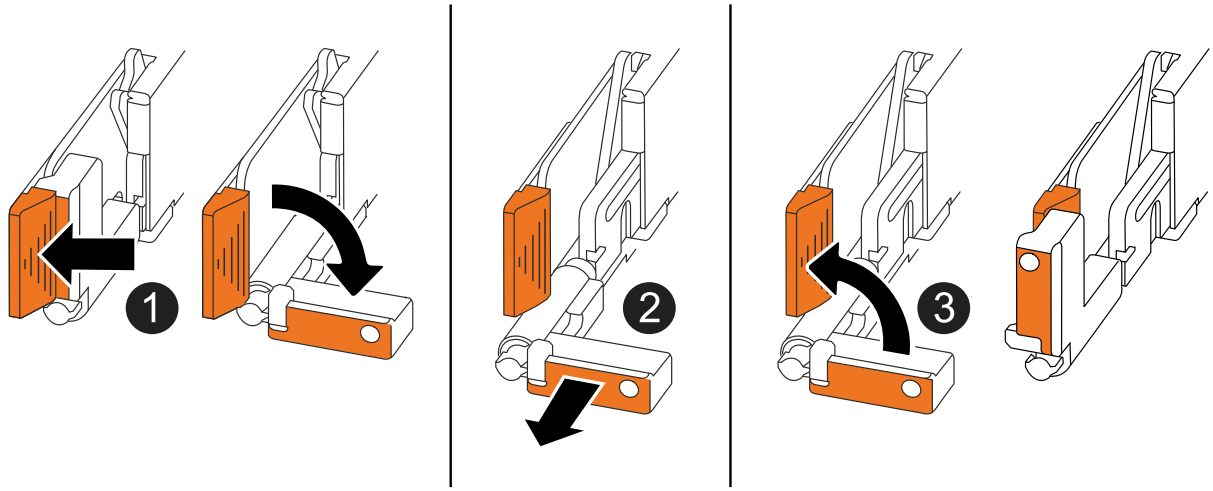

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace fan

To replace a fan, remove the failed fan and replace it with a new fan.

Steps

1. Identify the fan that you must replace by checking the console error messages.
2. Remove the failed fan:

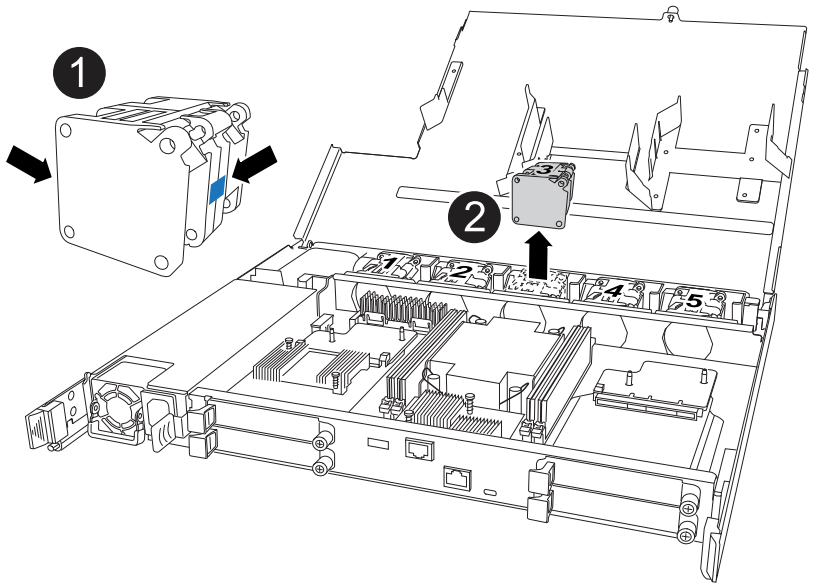

1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

3. Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

Step 4: Reinstall the controller module

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

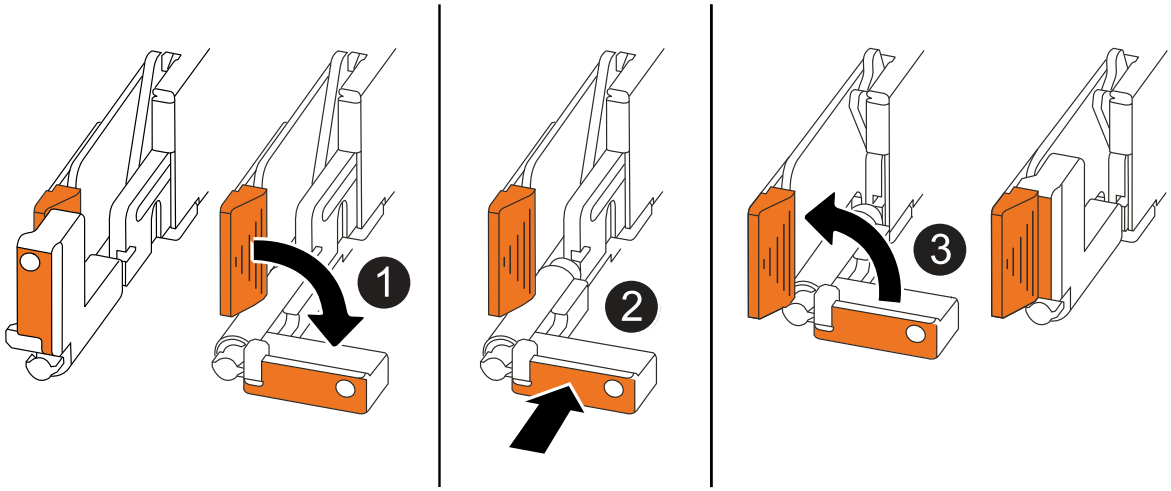

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Overview of I/O module maintenance - AFF C30 and AFF C60

The AFF C30 and AFF C60 storage systems offer flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding, hot-swapping, or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your storage system with the same type of I/O module, or with a different type of I/O module. You can hot-swap a cluster and HA I/O module when your storage system meets specific requirements. You can also add an I/O module to a storage system with available slots.

- Add an I/O module

  Adding additional I/O modules can improve redundancy, helping to ensure that the storage system remains operational even if one I/O module fails.

- Hot-swap a cluster and HA I/O module

Hot-swapping a failed cluster and HA I/O module can restore the storage system to its optimal operating state. Hot-swapping is done without having to manually take over the impaired controller.

To use this procedure, your storage system must be running ONTAP 9.17.1 or later and meet specific system requirements.

- Replace an I/O module

  Replacing a failing I/O module can restore the storage system to its optimal operating state.

Add an I/O module - AFF C30 and AFF C60

Add an I/O module to your AFF C30 or AFF C60 storage system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your AFF C30 and AFF C60 storage systems when there are slots available or when all slots are fully populated.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller module

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, remove the I/O blanking module from the target slot.

   Unused I/O slots should have blanking module installed to prevent possible thermal issues and for EMC compliance.

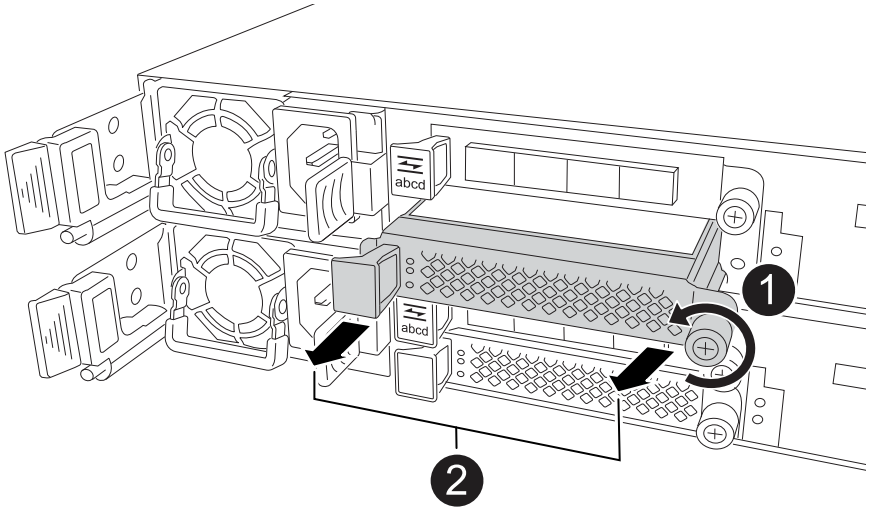

1	On the I/O blanking module, turn the thumbscrew counterclockwise to loosen.
2	Pull the I/O blanking module out of the controller using the tab on the left and the thumbscrew.

3. Install the new I/O module:
   a. Align the I/O module with the edges of the controller slot opening.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O module.

   c. Turn the thumbscrew clockwise to tighten.
4. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.

5. Reboot the impaired controller from the LOADER prompt: `bye`

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

6. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name.
```

7. Repeat these steps to add an I/O module to the other controller.
8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation: +

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, unplug any cabling on the target I/O module.
3. Remove the target I/O module from the controller:

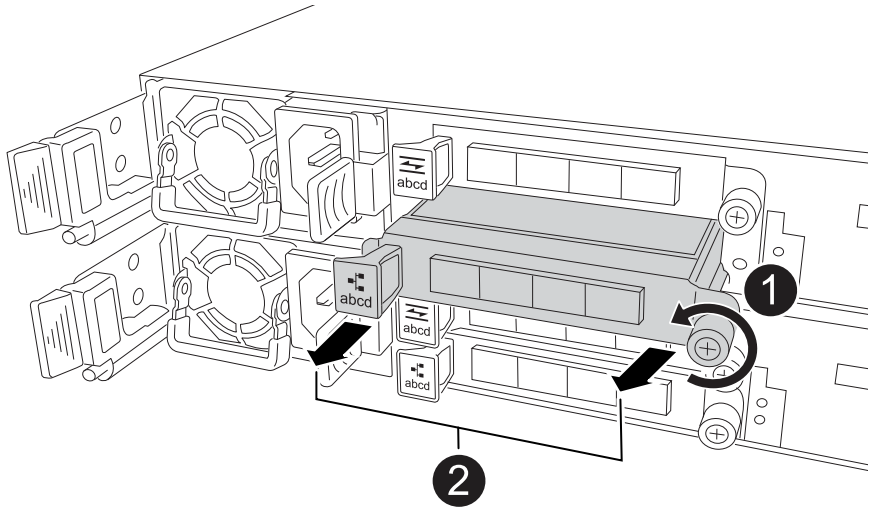

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the new I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
5. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
6. Repeat the I/O module remove and install steps to add any additional I/O modules in the controller.
7. Reboot the impaired controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.
8. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```
9. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```
10. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```
11. If you installed a NIC module, specify the usage mode for each port as *network*:

    ```
    storage port modify -node node_name -port port_name -mode network
    ```
12. Repeat these steps for the other controller.

Hot-swap the I/O module used for cluster and HA traffic - AFF C30 and AFF C60

The cluster and HA I/O module supports interconnects for clustering and high-availability. You can hot-swap the module in your AFF C30 or AFF C60 storage system when the module fails and if your storage system meets specific requirements.

To hot-swap a module, you ensure your storage system meets the procedure requirements, prepare the storage system and I/O module in slot 4, hot-swap the failed module for an equivalent one, bring the replacement module online, restore the storage system to normal operation, and return the failed module to NetApp.

About this task

- Hot-swapping the cluster and HA I/O module means that you do not have to perform a manual takeover; the impaired controller (the controller with the failed cluster and HA I/O module) has automatically taken over the healthy controller.

  When the impaired controller has taken over the healthy controller, the only way to recover without an outage is to hot-swap the module.

- It is critical to apply the commands to the correct controller when you are hot-swapping the cluster and HA I/O module:
  - The *impaired controller* is the controller on which you are hot-swapping the cluster and HA I/O module and it is the controller that has taken over the healthy controller.
  - The *healthy controller* is the HA partner of the impaired controller and it is the controller that was taken over by the impaired controller.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, make sure your storage system meets all requirements.

> If your storage system does not meet all requirements, you must use the replace an I/O module procedure.

- Your storage system must be running ONTAP 9.17.1 or later.
- The I/O module that failed must be a cluster and HA I/O module in slot 4 and you must be replacing it with an equivalent cluster and HA I/O module. You cannot change the I/O module type.
- Your storage system configuration must have only one cluster and HA I/O module located in slot 4, not two cluster and HA I/O modules.
- Your storage system must be a two-node (switchless or switched) cluster configuration.
- The controller with the failed cluster and HA I/O module (the impaired controller) must have already taken over the healthy partner controller. The takeover should have occurred automatically if the I/O module is failed.

  For two-node clusters, the storage system cannot discern which controller has the failed I/O module, so either controller might initiate the takeover. The cluster and HA I/O module hot-swap procedure is only supported when the controller with the failed I/O module (the impaired controller) has taken over the healthy controller.

You can verify that the impaired controller successfully took over the healthy controller by entering the `storage failover show` command.

If you are not sure which controller has the failed I/O module, contact NetApp Support.

- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Step 2: Prepare the storage system and I/O module slot 4

Prepare the storage system and I/O module slot 4 so that it is safe to remove the failed cluster and HA I/O module:

Steps

1. Properly ground yourself.
2. Unplug cabling from the failed cluster and HA I/O module.

   Make sure to label the cables so that later in this procedure you can reconnect them to the same ports.

3. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<number of hours down>h`

   For example, the following AutoSupport message suppresses automatic case creation for two hours:

   `node2::> system node autosupport invoke -node * -type all -message MAINT=2h`

4. Disable automatic giveback:
   a. Enter the following command from the console of the impaired controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
5. Prepare the failed cluster and HA module in slot 4 for removal by removing it from service and powering it off:
   a. Enter the following command:

      `system controller slot module remove -node impaired_node_name -slot slot_number`

   b. Enter `y` when you see the prompt *Do you want to continue?*

      For example, the following command prepares the module in slot 4 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

```
node2::> system controller slot module remove -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered off for removal.

Do you want to continue? {y|n}: y

The module has been successfully removed from service and powered
off. It can now be safely removed.
```

6. Verify the failed cluster and HA module in slot 4 is powered off:

   `system controller slot module show`

   The output should show *`powered-off`* in the status column for the failed module in slot 4.

Step 3: Replace the failed cluster and HA I/O module

Replace the failed cluster and HA I/O module in slot 4 with an equivalent I/O module:

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the failed cluster and HA I/O module from the impaired controller:

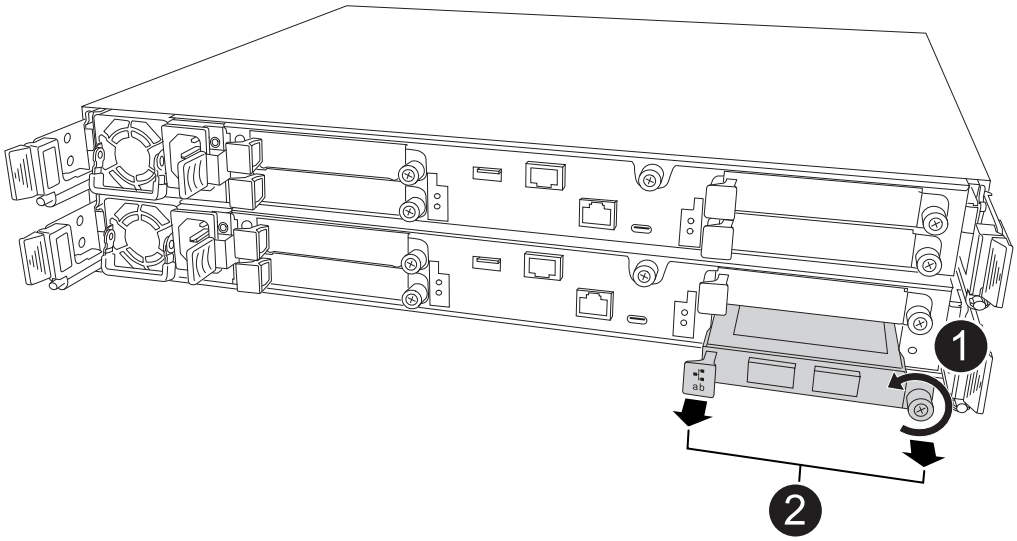

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew on the right.

3. Install the replacement cluster and HA I/O module into slot 4:

a. Align the I/O module with the edges of the slot.

b. Gently push the I/O module all the way into the slot, making sure to properly seat the I/O module into the connector.

   You can use the tab on the left and the thumbscrew on the right to push in the I/O module.

c. Turn the thumbscrew clockwise to tighten.

4. Cable the cluster and HA I/O module.

Step 4: Bring the replacement cluster and HA I/O module online

Bring the replacement cluster and HA I/O module in slot 4 online, verify the module ports initialized successfully, verify slot 4 is powered on, and then verify the module is online and recognized.

Steps

1. Bring the replacement cluster and HA I/O module online:

   a. Enter the following command:

      `system controller slot module insert -node` *`impaired_node_name`* `-slot` *`slot_name`*

   b. Enter `y` when you see the prompt, *Do you want to continue?*

      The output should confirm the cluster and HA I/O module was successfully brought online (powered on, initialized, and placed into service).

      For example, the following command brings slot 4 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered on and initialized.

Do you want to continue? {y|n}: `y`

The module has been successfully powered on, initialized and placed
into service.
```

2. Verify that each port on the cluster and HA I/O module successfully initialized:

   `event log show -event *hotplug.init*`

   > It might take several minutes to allow for any required firmware updates and port initialization.

   The output should show a hotplug.init.success EMS event logged for each port on the cluster and HA I/O module with *`hotplug.init.success:`* in the *`Event`* column.

For example, the following output shows initialization succeeded for cluster and HA I/O module ports e4b and e4a:

```
node2::> event log show -event *hotplug.init*

Time                Node             Severity      Event

------------------- ---------------- -------------
---------------------------

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4b" in slot 4 succeeded

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4a" in slot 4 succeeded

2 entries were displayed.
```

3. Verify I/O module slot 4 is powered on and ready for operation:

   `system controller slot module show`

   The output should show slot 4 status as *`powered-on`* and therefore ready for operation of the replacement cluster and HA I/O module.

4. Verify that the replacement cluster and HA I/O module is online and recognized.

   Enter the command from the console of the impaired controller:

   `system controller config show -node local -slot4`

   If the replacement cluster and HA I/O module was successfully brought online and is recognize, the output shows I/O module information, including port information, for slot 4.

   For example, you should see output similar to the following:

```
node2::> system controller config show -node local -slot 4

Node: node2
Sub- Device/
Slot slot Information
---- ---- ------------------------------
   4    - Dual 40G/100G Ethernet Controller CX6-DX
                  e4a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2807GJFM-B
                  e4b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2809G26F-A
                  Device Type:        CX6-DX PSID(NAP0000000027)
                  Firmware Version:   22.44.1700
                  Part Number:        111-05341
                  Hardware Revision:  20
                  Serial Number:      032403001370
```

Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the healthy controller, restoring automatic giveback, and reenabling AutoSupport automatic case creation.

Steps

1. Return the healthy controller (the controller that was taken over) to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode healthy_node_name
   ```

2. Restore automatic giveback from the console of the impaired controller (the controller that took over the healthy controller):

   ```
   storage failover modify -node local -auto-giveback true
   ```

3. If AutoSupport is enabled, restore automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=end
   ```

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an I/O module - AFF C30 and AFF C60

Replace an I/O module in your AFF C30 or AFF C60 storage system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

Before you begin

All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace a failed I/O module

To replace a failed I/O module, locate it in the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug cabling from the failed I/O module.

   Make sure to label the cables so that you know where they came from.

3. Remove the failed I/O module from the controller:

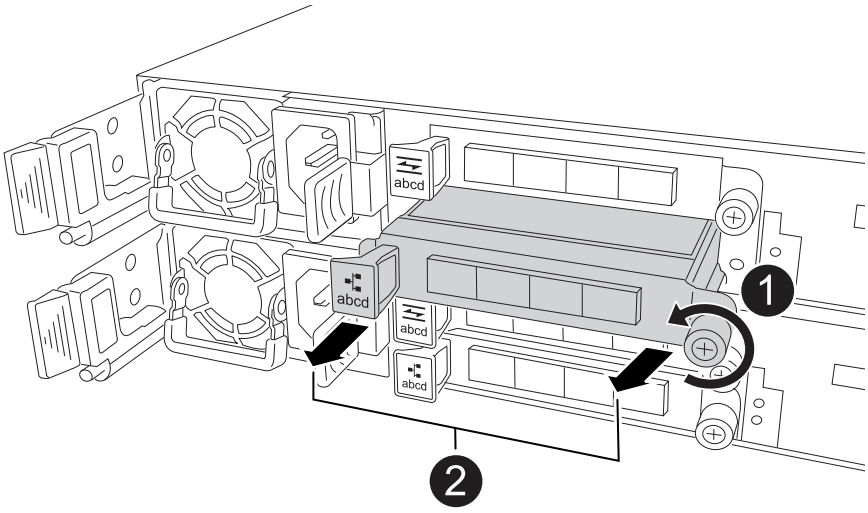

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the replacement I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

c. Turn the thumbscrew clockwise to tighten.

5. Cable the I/O module.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt:

```
bye
```

Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

3. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

4. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NV battery - AFF C30 and AFF C60

Replace the NV battery in your AFF C30 or AFF C60 storage system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

Before you begin

All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

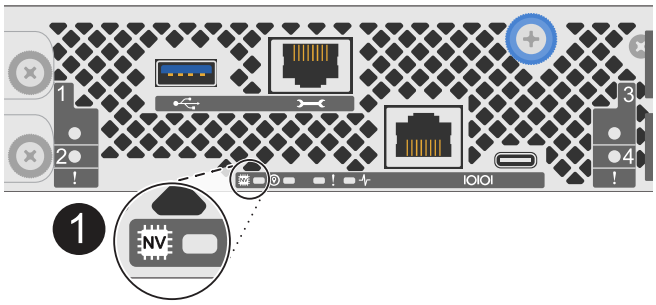

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

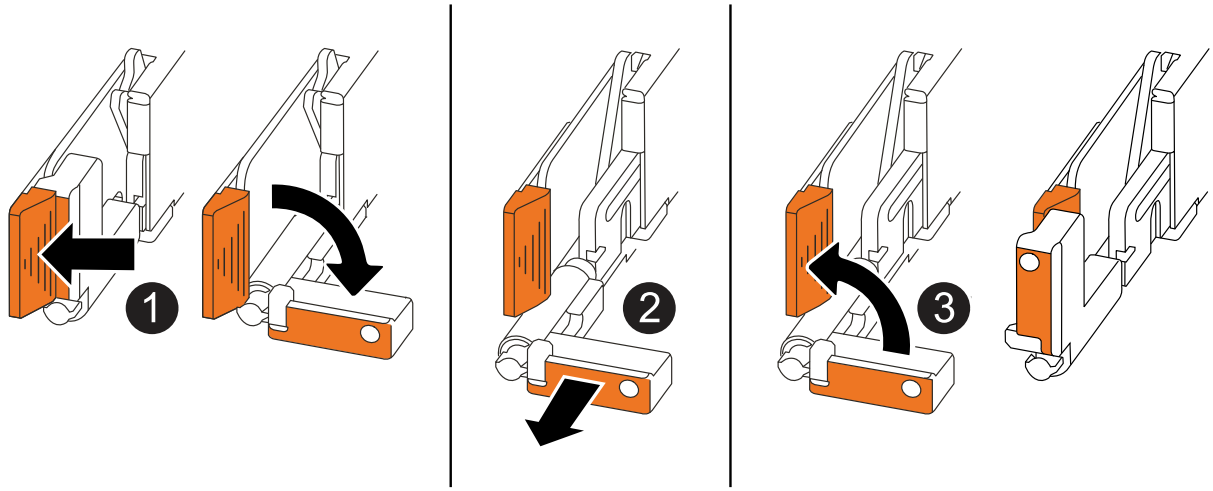

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the NV battery

Remove the failed NV battery from the controller and install the replacement NV battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the NV battery.
3. Remove the NV battery:

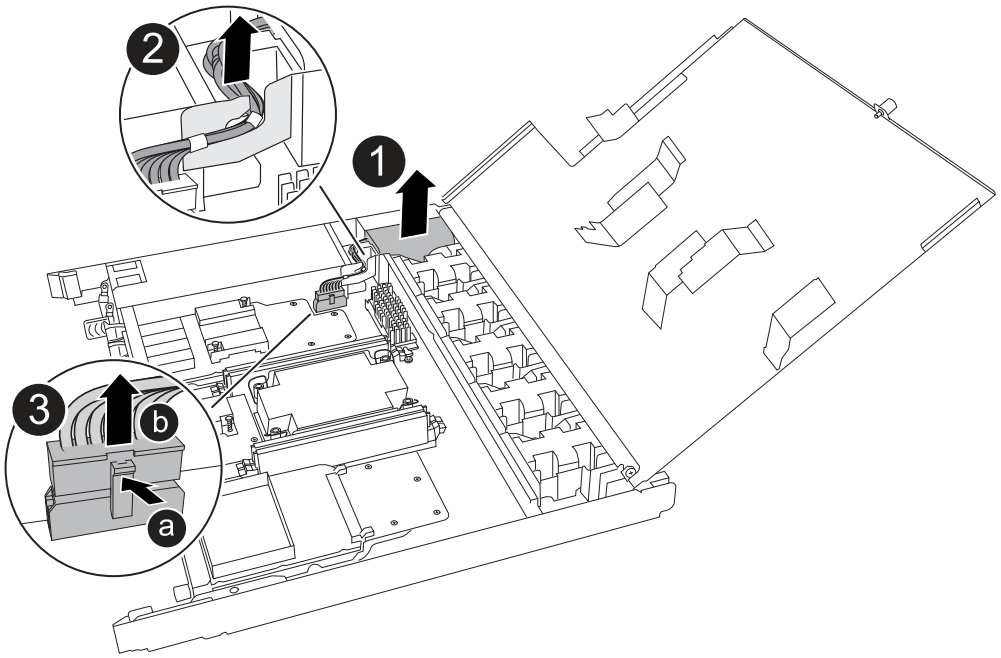

1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

4. Install the replacement NV battery:
   a. Remove the replacement battery from its package.
   b. Plug the wiring connector into its socket.
   c. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   d. Place the NV battery into its compartment.

      The NV battery should sit flush in its compartment.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

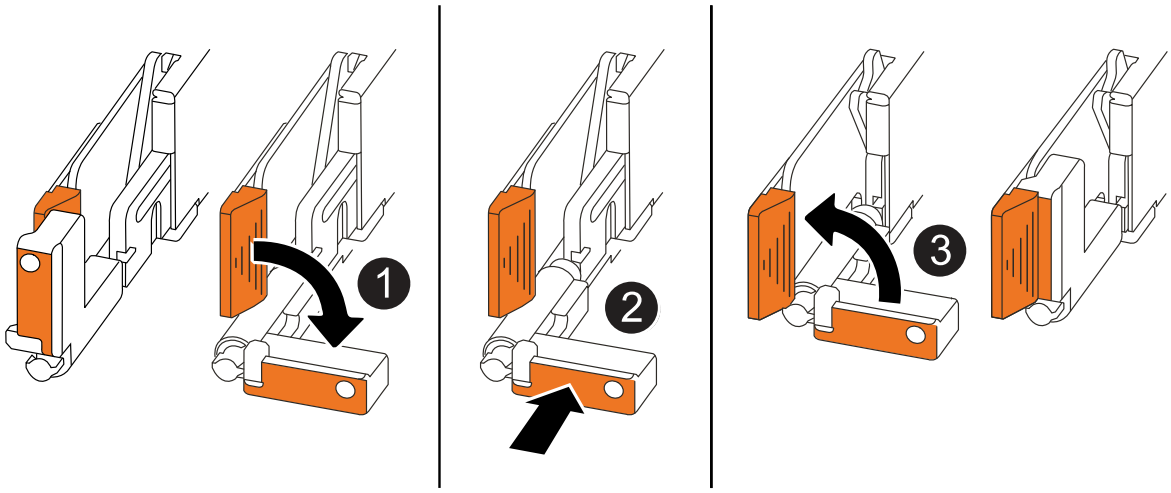

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.

The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - AFF C30 and AFF C60

Replace an AC or DC power supply unit (PSU) in your AFF C30 or AFF C60 storage system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cord, replacing the faulty PSU, and then reconnecting it to the power source.

About this task

- This procedure is written for replacing one PSU at a time.

The PSUs are redundant and hot-swappable.

- IMPORTANT: Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.
- Use the appropriate procedure for your type of PSU: AC or DC.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the power cord from the PSU by opening the power cord retainer, and then unplug the power cord from the PSU.

   PSUs do not have a power switch.

4. Remove the PSU:

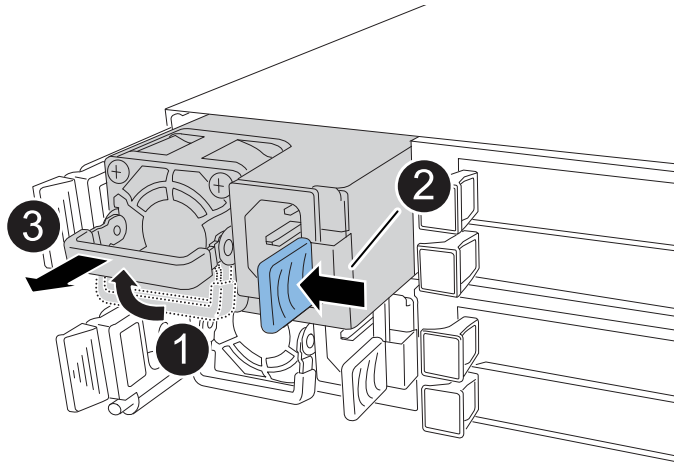

1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

5. Install the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

c. Rotate the handle down, so it is out of the way of normal operations.

6. Reconnect the power cord to the PSU and secure the power cord with the power cord retainer.

   Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the PSU:

   PSUs do not have a power switch.

   a. Unscrew the two thumb screws on the D-SUB DC power cord connector.

      The illustration and table in step 4 shows the two thumb screws (item #1) and the D-SUB DC power cord connector (item #2).

   b. Unplug the cord from the PSU and set it aside.

4. Remove the PSU:

   a. Rotate the handle up, to its horizontal position, and then grasp it.

   b. With your thumb, press the terracotta tab to release the locking mechanism.

   c. Pull the PSU out of the controller while using your other hand to support its weight.

   The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.

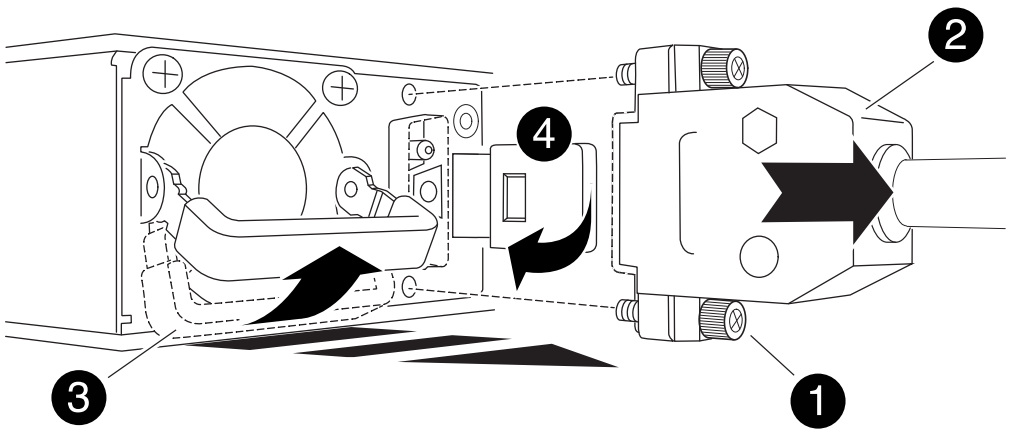

1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

5. Insert the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently slide the PSU into the controller until the locking tab clicks into place.

      A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.
6. Reconnect the D-SUB DC power cord:

   Once power is restored to the PSU, the status LED should be green.

   a. Plug the D-SUB DC power cord connector into the PSU.
   b. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF C30 and AFF C60

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your AFF C30 or AFF C60 storage system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

About this task

- You can use this procedure with all versions of ONTAP supported by your storage system.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

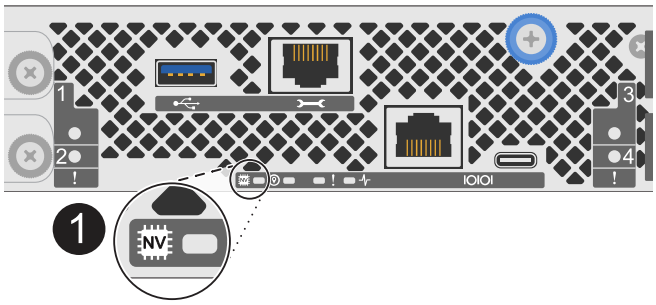

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

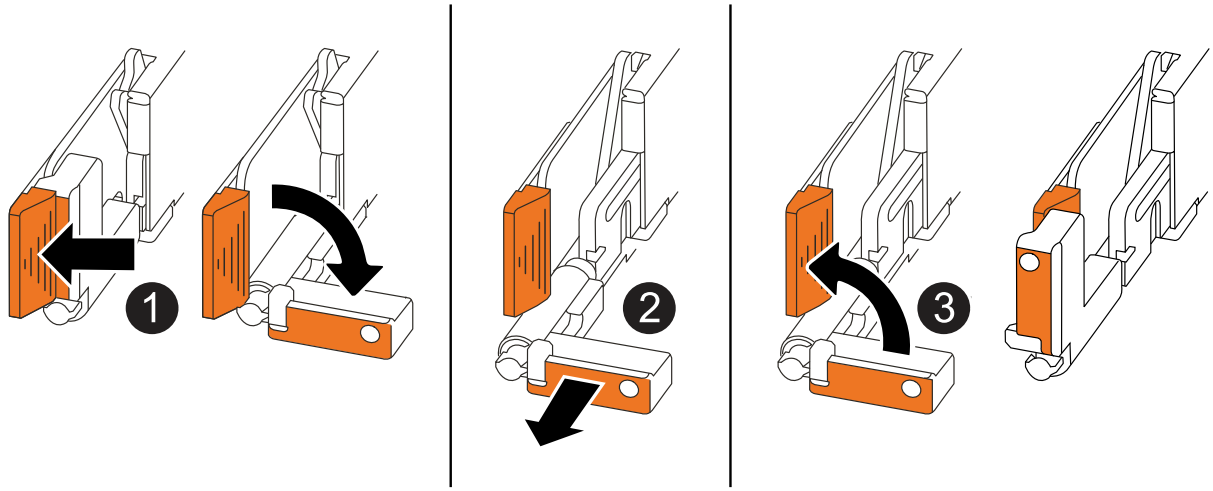

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the RTC battery

Remove the failed RTC battery and install the replacement RTC battery.

Steps

1. Locate the RTC battery.
2. Remove the RTC battery:

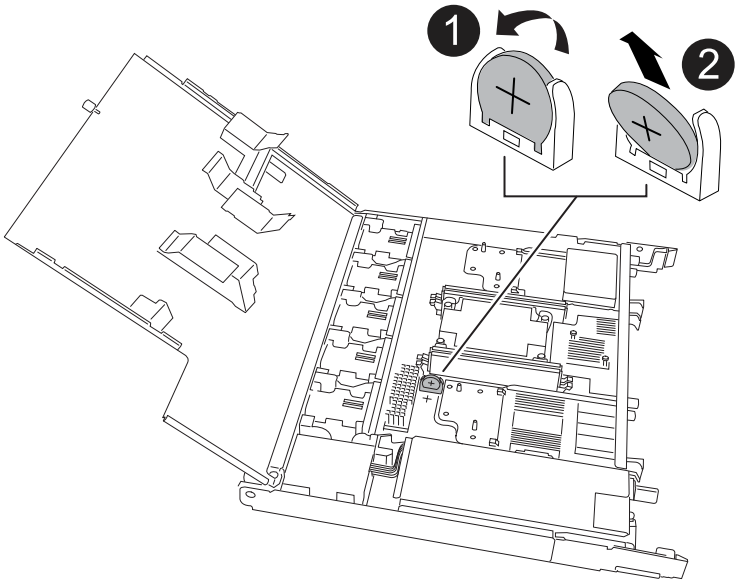

1	Gently rotate the RTC battery at an angle away from its holder.
2	Lift the RTC battery out of its holder.

3. Install the replacement RTC battery:
   a. Remove the replacement battery from the antistatic shipping bag.
   b. Position the battery so that the plus sign on the battery faces out to correspond with the plus sign on the motherboard.
   c. Insert the battery into the holder at an angle, and then push it into an upright position so it is fully seated in the holder.
   d. Visually inspect the battery to make sure that it is completely seated in its holder and that the polarity is correct.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

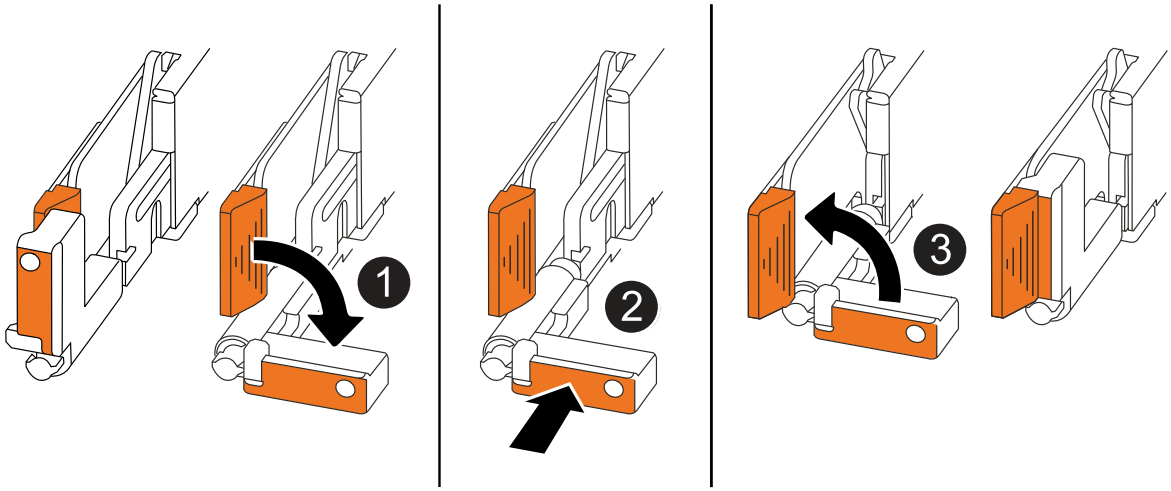

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Reset the time and date on the controller

After replacing the RTC battery, inserting the controller and powering on first BIOS reset, you will see the following error messages:
`RTC date/time error. Reset date/time to default`
`RTC power failure error`
These messages are expected and you can continue with this procedure.

1. On the healthy controller, check the date and time:

   `cluster date show`

   If your storage system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*.

2. On the impaired controller, at the LOADER prompt, check the time and date:

   `cluster date show`

   a. If necessary, modify the date:

      `set date mm/dd/yyyy`

   b. If necessary, set the time, in GMT:

```
set time hh:mm:ss
```

c. Confirm the date and time.

3. At the LOADER prompt, enter `bye` to reinitialize the I/O modules, other components, and let the controller reboot.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF C80 systems

Install and setup

Installation and configuration workflow - AFF C80

To install and configure your AFF C80 system, you review the hardware requirements, prepare your site, install and cable the hardware components, power on the system, and set up your ONTAP cluster.

Review installation requirements

Review the equipment and tools needed to install your storage system and storage shelves and review the lifting and safety precautions.

Prepare to install the AFF C80 storage system

To prepare to install your system, you need to get the site ready, check the environmental and electrical requirements, and ensure there's enough rack space. Then, unpack the equipment, compare its contents to the packing slip, and register the hardware to access support benefits.

Install the hardware for the AFF C80 storage system

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system in the cabinet or telco rack. Next, slide the shelves onto the rails. Finally, attach cable management devices to the rear of the storage system for organized cable routing.

Cable the controllers and storage shelves for AFF C80 storage system

To cable the hardware, first connect the storage controllers to your network and then connect the controllers to your storage shelves.

Power on the AFF C80 storage system

Before you power on the controllers, power on each NS224 shelf and assign a unique shelf ID to ensure each shelf is uniquely identified within the setup, connect the laptop or console to the controller, and then connect the controllers to the power sources.

6 Set up your cluster

After you've powered on your storage system, you set up your cluster.

Installation requirements - AFF C80

Review the equipment needed and the lifting precautions for your AFF C80 storage system and storage shelves.

Equipment needed for install

To install your storage system, you need the following equipment and tools.

- Access to a Web browser to configure your storage system
- Electrostatic discharge (ESD) strap
- Flashlight
- Laptop or console with a USB/serial connection
- Phillips #2 screwdriver

Lifting precautions

Storage systems and shelves are heavy. Exercise caution when lifting and moving these items.

Storage system weight

Take the necessary precautions when moving or lifting your storage system.

An A1K storage system can weigh up to 62.83 lbs (28.5 kg). To lift the storage system, use two people or a hydraulic lift.

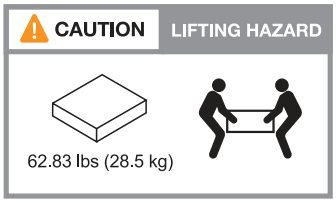

The storage system can weigh up to 151.68 lbs (68.8 kg). To lift the storage system, use four people or a hydraulic lift.

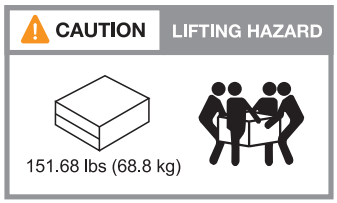

Shelf weight

Take the necessary precautions when moving or lifting your shelf.

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep

all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

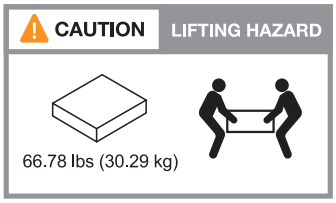

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

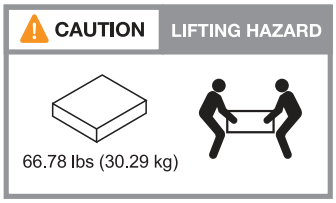

Related information

- Safety information and regulatory notices

What's next?

After you've reviewed the hardware requirements, you prepare to install your AFF C80 storage system.

Prepare to install - AFF C80

Prepare to install your AFF C80 storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the system to access support benefits.

Step 1: Prepare the site

To install your storage system, ensure that the site and the cabinet or rack that you plan to use meet specifications for your configuration.

Steps

1. Use NetApp Hardware Universe to confirm that your site meets the environmental and electrical requirements for your storage system.
2. Make sure you have adequate cabinet or rack space for your storage system, shelves, and any switches:
   - 4U in an HA configuration
   - 2U for each NS224 storage shelf
3. Install any required network switches.

   See the Switch documentation for installation instructions and NetApp Hardware Universe for compatibility information.

Step 2: Unpack the boxes

After you've ensured that the site and the cabinet or rack that you plan to use for your storage system meet the required specifications, unpack all boxes and compare the contents to the items on the packing slip.

Steps

1. Carefully open all the boxes and lay out the contents in an organized manner.
2. Compare the contents you've unpacked with the list on the packing slip.

   You can get your packing list by scanning the QR code on the side of the shipping carton.

   The following items are some of the contents you might see in the boxes.

   Ensure that everything in the boxes matches the list on the packing slip. If there are any discrepancies, note them down for further action.

Hardware

- Bezel
- Cable management device
- Storage system
- Rail kits with instructions (optional)
- Storage shelf (if you ordered additional storage)

Cables

- Management Ethernet cables (RJ-45 cables)
- Network cables
- Power cords
- Storage cables (if you ordered additional storage)
- USB-C serial console cable

Step 3: Register your storage system

After you've ensured that your site meets the requirements for your storage system specifications, and you've verified that you have all the parts you ordered, you should register your storage system.

Steps

1. Locate the System Serial Numbers (SSN) for every controller being installed. You can find the serial numbers in the following locations:
2. You can find the serial numbers in the following locations:
   - On the packing slip
   - In your confirmation email
   - On each controller's System Management module

3. Go to the NetApp Support Site.
4. Determine whether you need to register your storage system:

If you are a…	Follow these steps…
Existing NetApp customer	a. Sign in with your username and password. b. Select **Systems** > **My Systems**. c. Confirm that the new serial numbers are listed. d. If it is not, follow the instructions for new NetApp customers.
New NetApp customer	a. Click **Register Now**, and create an account. b. Select **Systems** > **Register Systems**. c. Enter the storage system's serial numbers and requested details. After your registration is approved, you can download any required software. The approval process might take up to 24 hours.

What's next?

After you've prepared to install your AFF C80hardware, you install the hardware for your AFF C80 storage system.

Install the hardware - AFF C80

After you prepare to install your AFF C80 storage system, install the hardware for the system. First, install the rail kits. Then install and secure your storage system in a cabinet or telco rack.

Skip this step if your cabinet is pre-populated.

Before you begin

- Make sure you have the instructions packaged with the rail kit.
- Be aware of the safety concerns associated with the weight of the storage system and shelf.
- Understand that the airflow through the storage system enters from the front where the bezel or end caps are installed and exhausts out the rear where the ports are located.

Steps

1. Install the rail kits for your storage system and shelves as needed, using the instructions included with the kits.
2. Install and secure your storage system in the cabinet or telco rack:
   a. Position the storage system onto the rails in the middle of the cabinet or telco rack, and then support the storage system from the bottom and slide it into place.
   b. Make sure that the guiding pins of the cabinet or telco rack are securley in the chassis guide slots.
   c. Secure the storage system to the cabinet or telco rack using the included mounting screws.
3. Attach the bezel to the front of the storage system.
4. Attach the cable management devices to the rear of the storage system.

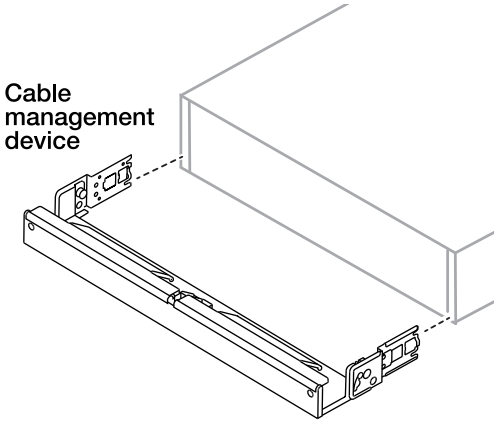

5. Install and secure the shelf as needed.
   a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

      If you are installing multiple shelves, place the first shelf directly above the controllers. Place the second shelf directly under the controllers. Repeat this pattern for any additional shelves.

   b. Secure the shelf to the cabinet or telco rack using the included mounting screws.

      What's next?

      After you've installed the hardware for your AFF C80 system, you cable the hardware for your AFF C80 storage system.

Cable the hardware - AFF C80

After you install the rack hardware for your AFF C80 storage system, install the network cables for the controllers, and connect the cables between the controllers and storage shelves.

Before you begin

Contact your network administrator for information about connecting the storage system to the switches.

About this task

- These procedures show common configurations. The specific cabling depends on the components ordered for your storage system. For comprehensive configuration and slot priority details, see NetApp Hardware Universe.
- The I/O slots on an AFF C80 controller are numbered 1 through 11.

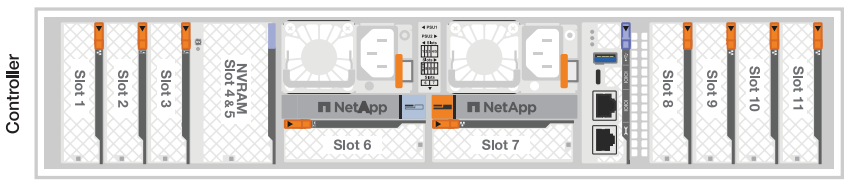

- The cabling graphics have arrow icons showing the proper orientation (up or down) of the cable connector pull-tab when inserting a connector into a port.

  As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it

over and try again.

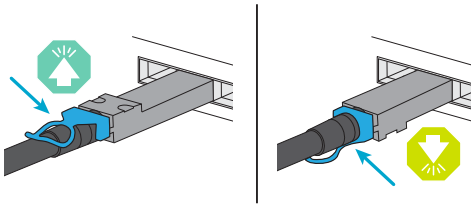

- If cabling to an optical switch, insert the optical transceiver into the controller port before cabling to the switch port.

Step 1: Cable the cluster/HA connections

Cable the controllers to your ONTAP cluster. This procedure differs depending on your storage system model and I/O module configuration.

The cluster interconnect traffic and the HA traffic share the same physical ports.

Switchless cluster cabling

Use the the Cluster/HA interconnect cable to connect to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A to port e1a on Controller B.
2. Connect port e7a on Controller A to port e7a on Controller B.

Cluster/HA interconnect cables

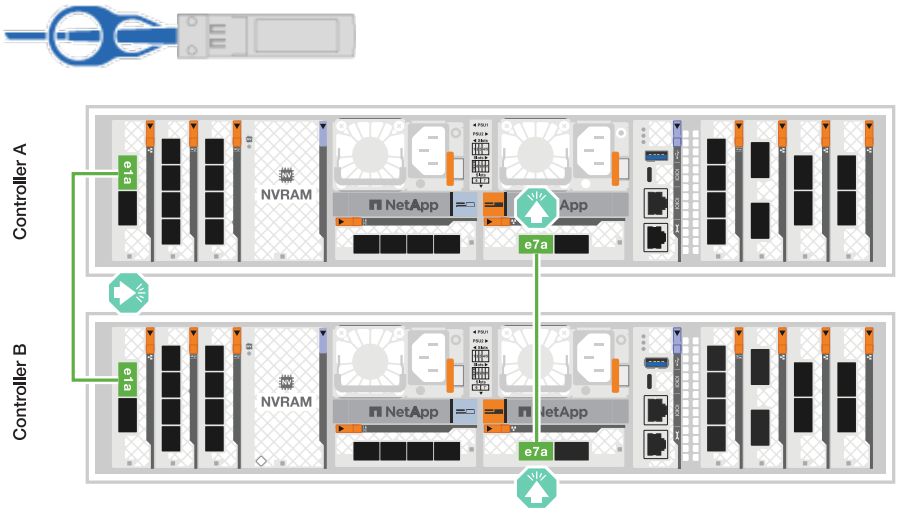

Switched cluster cabling

Use the 100 GbE cable to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A and port e1a on Controller B to cluster network switch A.
2. Connect port e7a on Controller A and port e7a on Controller B to cluster network switch B.

100 GbE cable

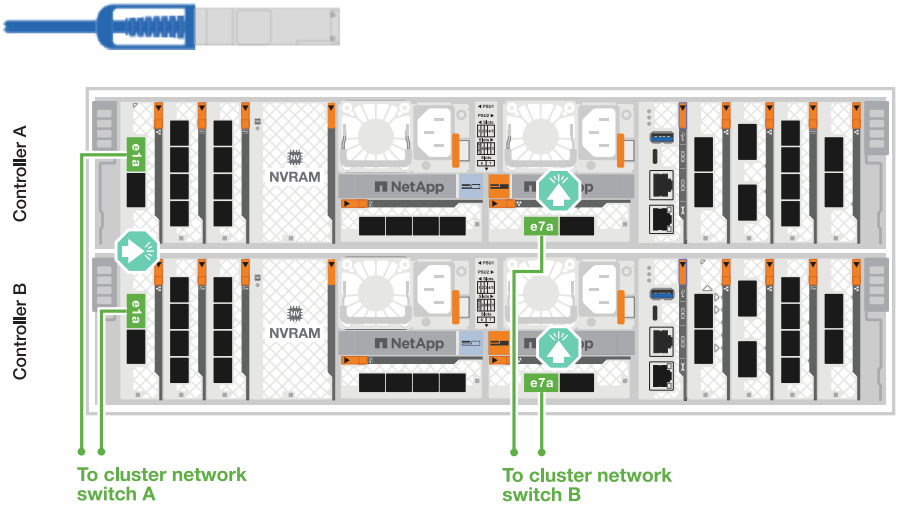

Step 2: Cable the host network connections

Connect the Ethernet module ports to your host network.

The following are some typical host network cabling examples. See NetApp Hardware Universe for your specific system configuration.

Steps

1. Connect ports e9a and e9b to your Ethernet data network switch.

   For maximum system performance for cluster and HA traffic, do not use ports e1b and e7b ports for host network connections. Use a separate host card to maximize performance.

   100 GbE cable

   

   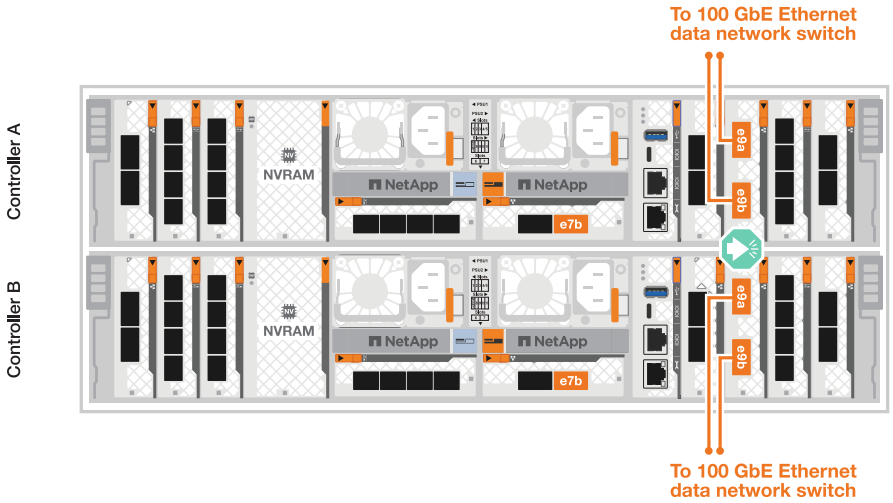

2. Connect your 10/25 GbE host network switches.

   4-ports, 10/25 GbE Host

   

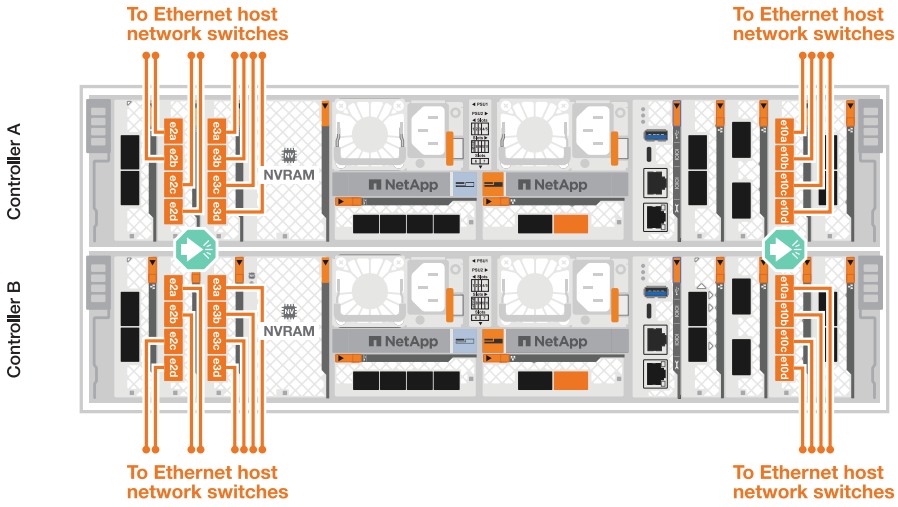

Step 3: Cable the management network connections

Use the 1000BASE-T RJ-45 cables to connect the management (wrench) ports on each controller to the management network switches.

1000BASE-T RJ-45 cables

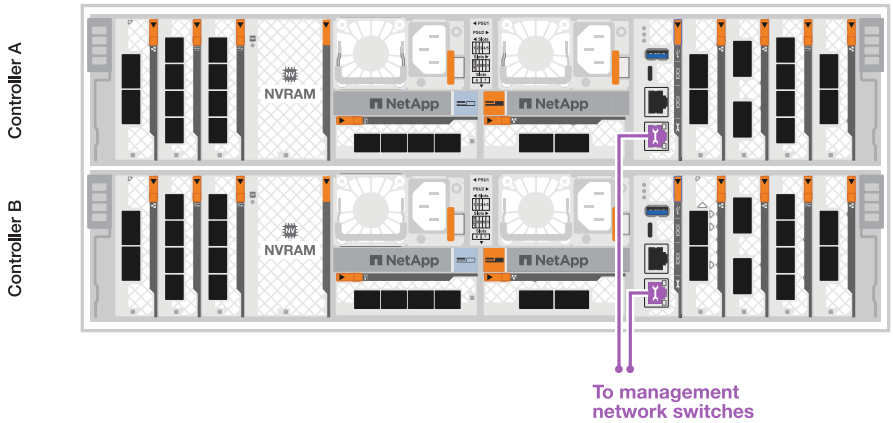

Do not plug in the power cords yet.

Step 4: Cable the shelf connections

The following cabling procedures show how to connect your controllers to a storage shelf. Choose one of the following cabling options that matches your setup.

For the maximum number of shelves supported for your storage system and for all of your cabling options, see NetApp Hardware Universe.

Option 1: One NS224 storage shelf

Connect each controller to the NSM modules on the NS224 shelf. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables

Steps

1. Connect controller A port e11a to NSM A port e0a.
2. Connect controller A port e11b to port NSM B port e0b.

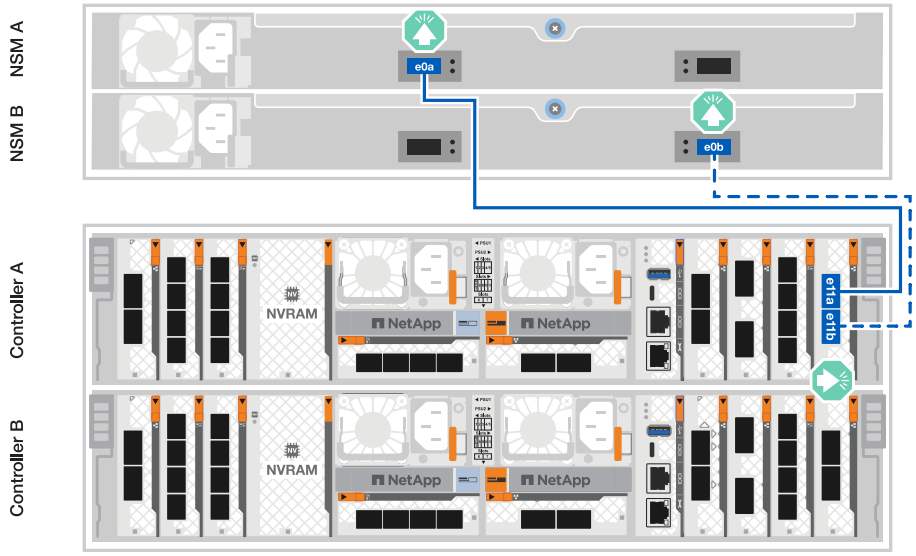

3. Connect controller B port e11a to NSM B port e0a.
4. Connect controller B port e11b to NSM A port e0b.

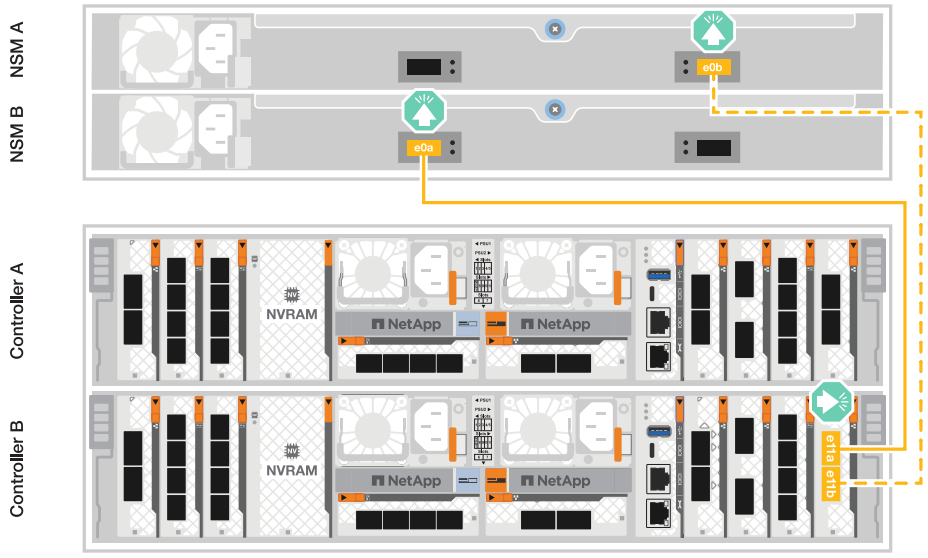

Option 2: Two NS224 storage shelves

Connect each controller to the NSM modules on both NS224 shelves. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables

Steps

1. On controller A, connect the following ports:
   a. Connect port e11a to shelf 1, NSM A port e0a.
   b. Connect port e11b to shelf 2, NSM B port e0b.
   c. Connect port e8a to shelf 2, NSM A port e0a.
   d. Connect port e8b to shelf 1, NSM B port e0b.

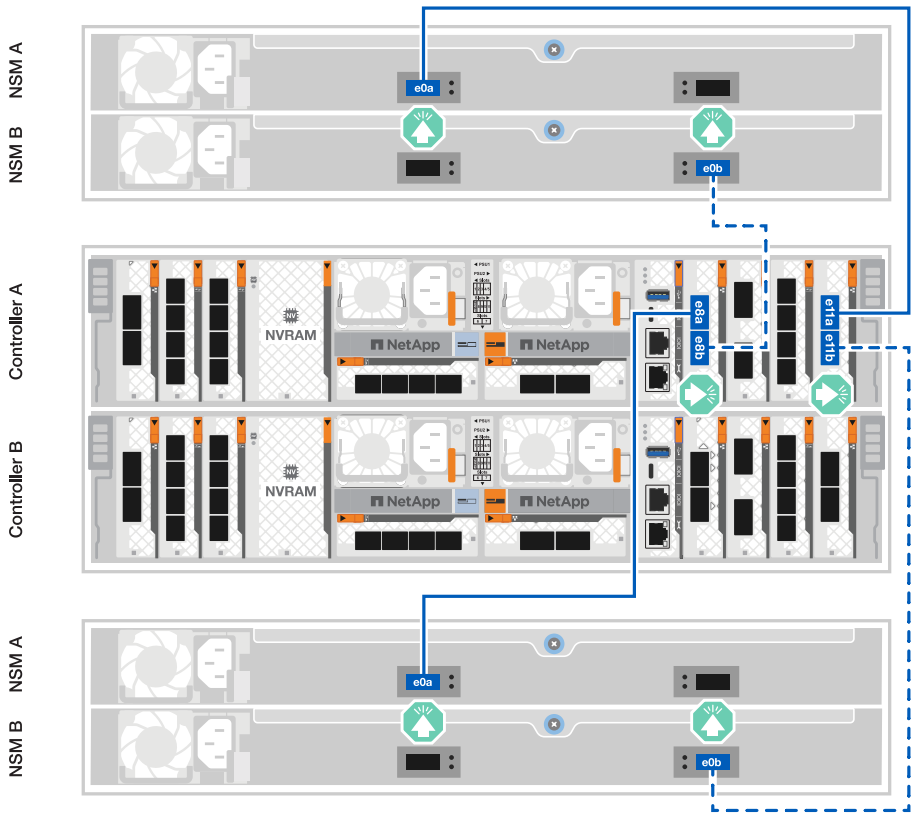

2. On controller B, connect the following ports:
   a. Connect port e11a to shelf 1, NSM B port e0a.
   b. Connect port e11b to shelf 2, NSM A port e0b.
   c. Connect port e8a to shelf 2, NSM B port e0a.
   d. Connect port e8b to shelf 1, NSM A port e0b.

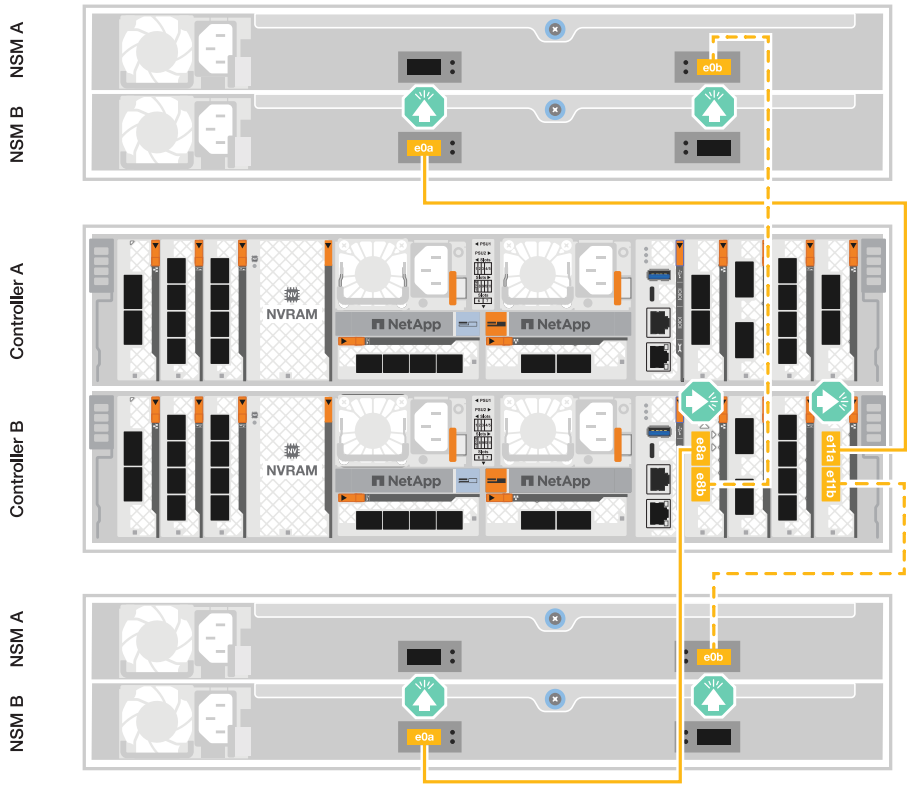

What's next?

After you've cabled the hardware for your AFF C80 system, you power on the AFF C80 storage system.

Power on the storage system - AFF C80

After you install the rack hardware for your AFF C80 storage system and install the cables for the controllers and storage shelves, you should power on your storage shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

Each shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup.

Before you begin

Make sure you have a paperclip or narrow tipped ball point pen for setting NS224 storage shelf IDs.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed shelf ID of 00.

- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, and then plug them back in) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, and then connecting the power cords to power sources on different circuits.

   The shelf powers on and boots automatically when plugged into the power source.

2. Remove the left end cap to access the shelf ID button behind the faceplate.

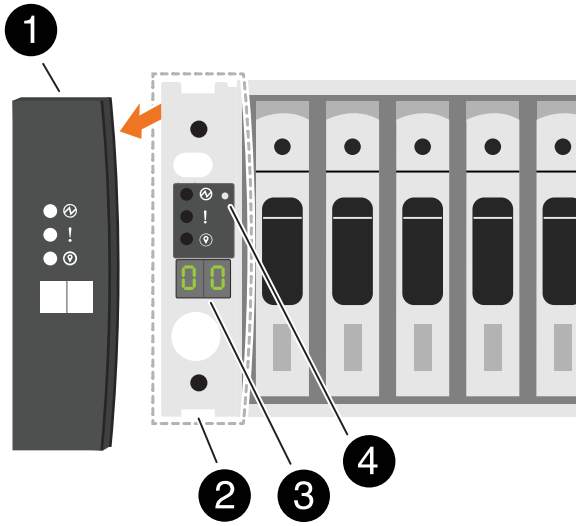

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID number
4	Shelf ID button

3. Change the first number of the shelf ID:
   a. Insert the straightened end of a paperclip or narrow tipped ball point pen into the small hole to press the shelf ID button.
   b. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

      It can take up to 15 seconds for the number to blink. This activates the shelf ID programming mode.

      If the ID takes longer than 15 seconds to blink, press and hold the shelf ID button again, making sure to press it in all the way.

   c. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

      Each press and release duration can be as short as one second.

      The first number continues to blink.

4. Change the second number of the shelf ID:
   a. Press and hold the button until the second number on the digital display blinks.

      It can take up to three seconds for the number to blink.

      The first number on the digital display stops blinking.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

      The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the shelf ID button until the second number stops blinking.

   It can take up to three seconds for the number to stop blinking.

   Both numbers on the digital display start blinking and the amber LED illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.
   a. Unplug the power cord from both power supplies on the shelf.
   b. Wait 10 seconds.
   c. Plug the power cords back into the shelf power supplies to complete the power cycle.

      A power supply is powered on as soon as the power cord is plugged in. Its bicolored LED should illuminate green.

7. Replace the left end cap.

Step 2: Power on the controllers

After you've powered on your shelves and assigned them unique IDs, power on the storage controllers.

Steps

1. Connect your laptop to the serial console port. This will allow you to monitor the boot sequence when the controllers are powered on.
   a. Set the serial console port on the laptop to 115,200 baud with N-8-1.

      See your laptop's online help for instructions on how to configure the serial console port.

   b. Using the console cable provided with your storage system, connect one end of the console cable to your laptop and the other end to the serial console port on controller A.
   c. Connect the laptop to the switch on the management subnet.

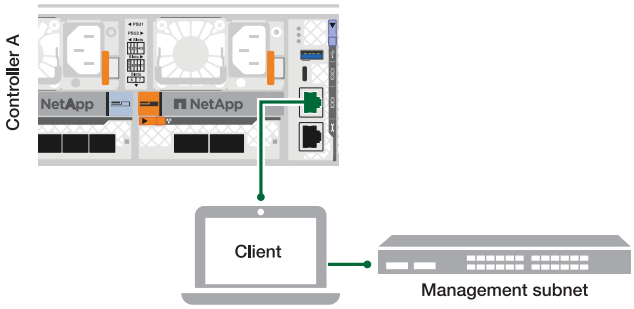

2. Assign a TCP/IP address to the laptop, using one that is on the management subnet.
3. Plug the two power cords into the controller power supplies, and then connect them to power sources on different circuits.

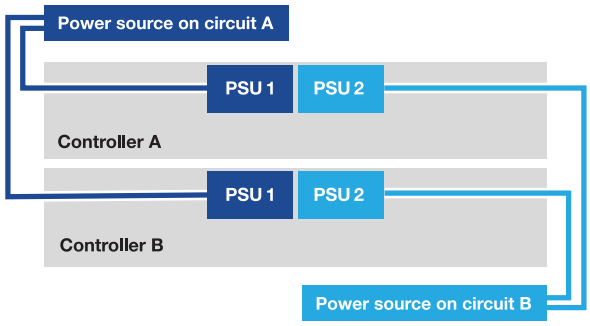

   - The system begins to boot. Initial booting might take up to eight minutes.
   - The LEDs flash on and the fans start, which indicates that the controllers are powering on.
   - The fans might be very noisy when they first start up. The fan noise during start-up is normal.
4. Secure the power cords using the securing device on each power supply.

What's next?

After you've turned on your AFF C80 storage system, you set up your cluster.

Maintain

Overview of the maintenance procedures - AFF C80

Maintain the hardware of your AFF C80 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the AFF C80 system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF C80 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller.
NVRAM	The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module.
NV battery	The NV battery is responsible for providing power to the NVRAM module while data in-flight is being destaged to flash memory after a power loss.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.
System Management module	The System Management module provides the interface between the controller and a console or laptop for controller or system maintenance purposes. The System Management module contains the boot media and stores the system serial number (SSN).

Boot media automated recovery workflow - AFF C80

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your AFF C80 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the System Management module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - AFF C80

Before replacing the boot media in your AFF C80 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming the cluster ports on the impaired controller are working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Review the following requirements.

- You must replace the failed boot media with a replacement boot media you received from NetApp.

- The cluster ports are used to communicate between the two controllers during the automated boot recovery process. Make sure that the cluster ports on the impaired controller are working properly.
- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg
  - /cfcard/kmip/certs/client.crt
  - /cfcard/kmip/certs/client.key
  - /cfcard/kmip/certs/CA.pem
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - AFF C80

Shut down the impaired controller in your AFF C80 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - AFF C80

The boot media in your AFF C80 storage system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media in the System Management module, and then reinstalling the System Management module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the System Management module and is accessed by removing the module from the system.

Steps

1. Make sure NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged. If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

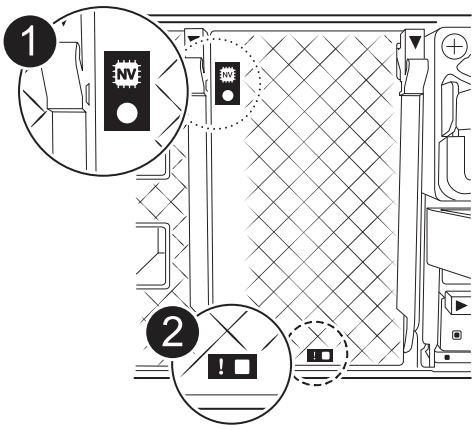

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller's PSUs.

   If your system has DC power, disconnect the power block from the PSUs.

   a. Remove any cables connected to the System Management module. Make sure to label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   c. Depress the system management cam button.
   The cam lever moves away from the chassis.
   d. Rotate the cam lever all the way down and remove the System Management module from the controller module.
   e. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:

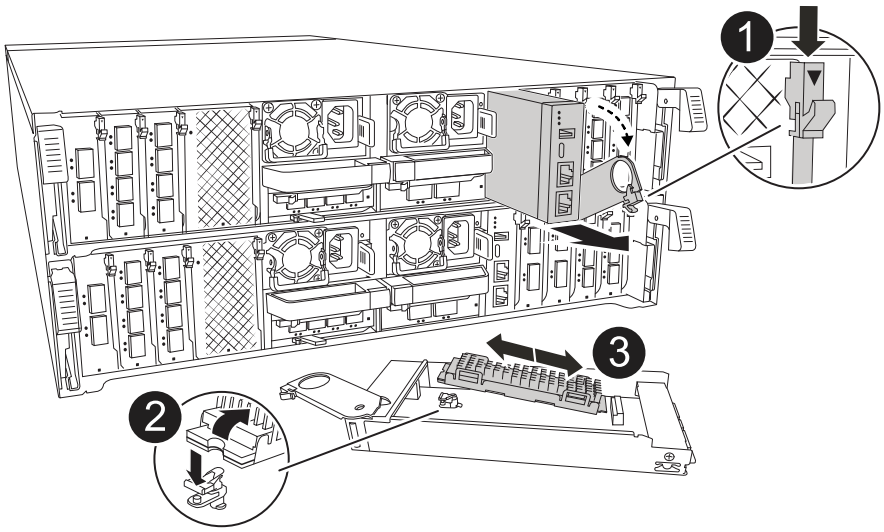

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue locking button.

b. Rotate the boot media up, slide it out of the socket, and set it aside.

5. Install the replacement boot media into the System Management module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down toward the locking button.

c. Push the locking button, rotate the boot media all the way down and then release the locking button.

6. Reinstall the System Management module:

a. Rotate the cable management tray up to the closed position.

b. Recable the System Management module.

7. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

If you have DC power supplies, reconnect the power block to the power supplies.

8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - AFF C80

After installing the new boot media device in your AFF C80 storage system, you can start the automated boot media recovery process to restore the configuration from the partner

node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

   Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

   Show example of configuration error finding prompts

   ```
   Error when fetching key manager config from partner ${partner_ip}:
   ${status}

   Has key manager been configured on this system

   Is the key manager onboard
   ```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media part to NetApp - AFF C80

If a component in your AFF C80 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF C80

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the AFF C80 system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

If your system is running in ONTAP 9.17.1 and later, use the automatic boot recovery procedure.

To get started, review the recovery requirements, shut down the controller, replace the boot media, use the USB drive to restore the image, and reapply encryption settings if necessary.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the impaired controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONTAP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - AFF C80

Before replacing the boot media in your AFF C80 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption support for manual boot media recovery - AFF C80

To ensure data security on your AFF C80 storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

What's next?

After checking the encryption key support and status on the boot media, you need to shut down the controller.

Shut down the controller for manual boot media recovery - AFF C80

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After shutting down the controller, you need to replace the boot media.

Replace the boot media and prepare for manual boot recovery - AFF C80

The boot media in your AFF C80 system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

Step 1: Replace the boot media

The boot media is located inside the System Management module and is accessed by removing the module from the system.

Steps

1. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
2. Unplug the controller's PSUs.

   If your system has DC power, disconnect the power block from the PSUs.

   a. Remove any cables connected to the System Management module. Make sure to label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.

   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

   c. Depress the system management cam button.
   The cam lever moves away from the chassis.

   d. Rotate the cam lever all the way down and remove the System Management module from the controller module.

   e. Place the System Management module on an anti-static mat, so that the boot media is accessible.
3. Remove the boot media from the management module:

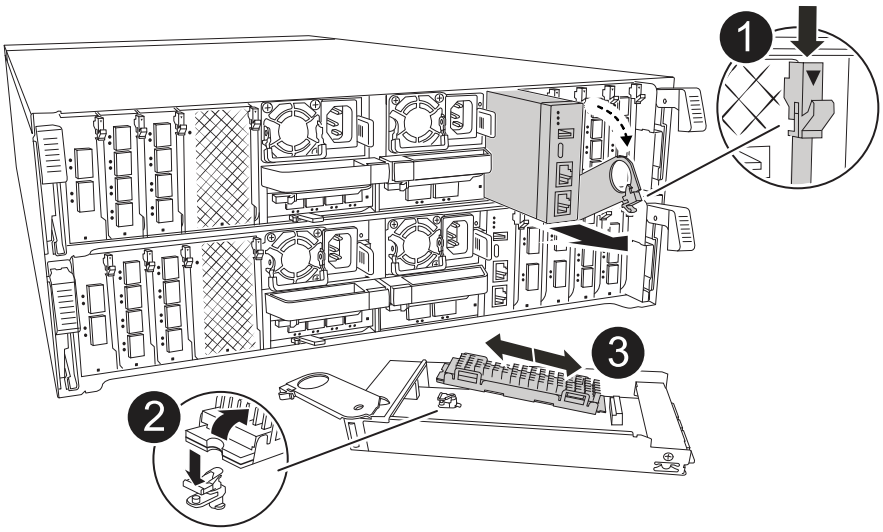

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue locking button.

b. Rotate the boot media up, slide it out of the socket, and set it aside.

4. Install the replacement boot media into the System Management module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down toward the locking button.

c. Push the locking button, rotate the boot media all the way down and then release the locking button.

5. Reinstall the System Management module:

a. Rotate the cable management tray up to the closed position.

b. Recable the System Management module.

Step 2: Transfer the boot image to the boot media

The replacement boot media that you installed is without an ONTAP image. You can transfer the ONTAP image to the replacement boot media by downloading the appropriate ONTAP service image from the NetApp Support Site to a USB flash drive and then to the replacement boot media.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- Download a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site. Use the `version -v` command to display if your version of ONTAP supports NVE. If the command output displays `<10no- DARE>`, your version of ONTAP does not support NVE.
  - If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption,

as indicated in the download button.

- If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image from the Downloads link on the page, to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      The USB flash drive should have the appropriate ONTAP image of what the impaired controller is running.

   c. Remove the USB flash drive from your laptop.
2. Insert the USB flash drive into the USB-A port on the System Management module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

3. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   If you have DC power supplies, reconnect the power block to the power supplies.

4. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

What's next?

After replacing the boot media, you need to boot the recovery image.

Manual boot media recovery from a USB drive - AFF C80

After installing the new boot media device in your AFF C80 system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

What's next?

After booting the recovery image, you need to restore encryption on the boot media.

Restore encryption keys after manual boot recovery - AFF C80

Restore encryption on the replacement boot media in your AFF C80 system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

What's next?

After restoring encryption on the boot media, you need to return the failed part to NetApp.

Return the failed part to NetApp - AFF C80

If a component in your AFF C80 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - AFF C80

Get started with replacing the chassis of your AFF C80 system by reviewing the replacement requirements, shutting down the controllers, replacing the chassis, and verifying system operations.

1 Review the chassis replace requirements

To replace the chassis, you must meet certain requirements.

2 Shut down the controllers

Shut down the controllers so you can perform maintenance on the chassis.

3 Replace the chassis

Replacing the chassis includes moving the power supplies, hard drives, and controller module from the impaired chassis to the new chassis, and swapping out the impaired chassis with the new chassis of the same model as the impaired chassis.

4 Complete chassis replacement

Verify the HA state of the chassis and return the failed part to NetApp.

Requirements to replace the chassis - AFF C80

Before replacing the chassis of your AFF C80 system, verify all other components in the system are functioning properly, ensure that you have local administrator credentials for ONTAP, the correct replacement chassis, and the necessary tools.

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Review the following requirements.

- Make sure all other components in the system are functioning properly; if not, contact NetApp support for assistance.
- Obtain local administrator credentials for ONTAP if you don't have them.
- Make sure that you have the necessary tools and equipment for the replacement.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

What's next?

After you've reviewed the requirements to replace the chassis, you need to shut down the controller.

Shut down the controller to replace the chassis - AFF C80

Shut down the controller in your AFF C80 storage system to prevent data loss and ensure system stability when replacing the chassis.

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

What's next?

After you've shut down the controller, you need to replace the chassis.

Replace the chassis - AFF C80

Replace the chassis of your AFF C80 system when a hardware failure requires it. The replacement process involves removing the controllers and power supply units (PSUs), removing the drives, installing the replacement chassis, and reinstalling the chassis components.

Step 1: Remove the PSUs and cables

You need to remove all four power supply units (PSUs), two per controller, before removing the controller. Removing them lightens the overall weight of each controller.

Steps

1. Remove the four PSUs:
   a. If you are not already grounded, properly ground yourself.
   b. Unplug power cords from the controller module PSU.

      If your system has DC power, disconnect the power block from the PSUs.

   c. Remove the PSU from the controller by rotating the PSU handle up so that you can pull the PSU out, press the PSU locking tab, and then pull PSU out of the controller module.

      The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

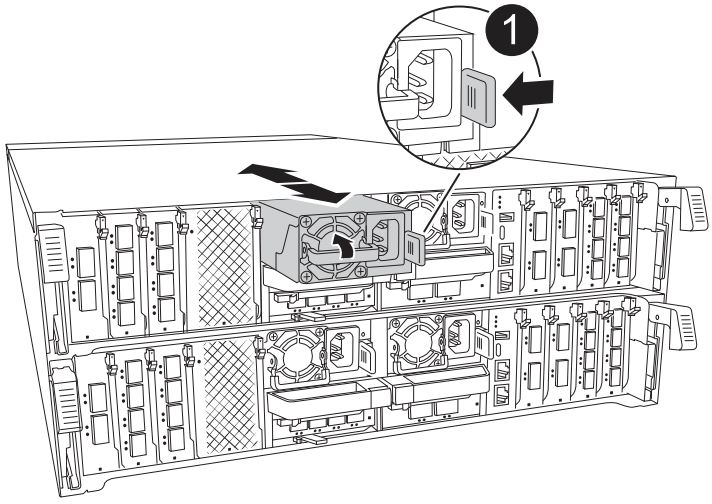

1	Terracotta PSU locking tab

d. Repeat these steps for the remaining PSUs.

2. Remove the cables:

a. Unplug the system cables and any SFP and QSFP modules (if needed) from the controller module, but leave them in the cable management device to keep them organized.

Cables should have been labeled at the beginning of this procedure.

b. Remove the cable management device from the controller modules and set them aside.

Step 2: Remove the controller modules and drives

Remove the controllers from the chassis and then remove the drives from the chassis.

Steps

1. Check the that amber NVRAM status LED located in slot 4/5 on the back of each controller module is off. Look for the NV icon.

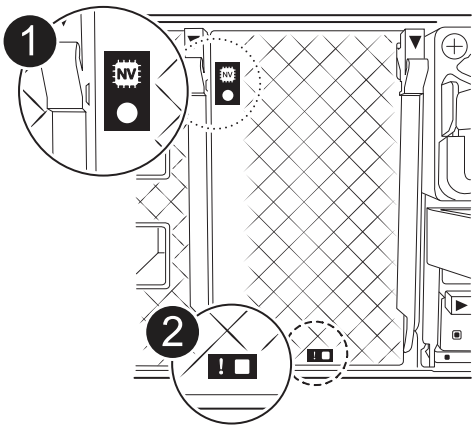

1	NVRAM status LED
2	NVRAM attention LED

◦ If the NVRAM LED is off, go to the next step.

◦ If the NVRAM LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact NetApp Support Site for assistance.

2. Remove the controller modules:

a. Press down on both of the locking latches on the controller, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

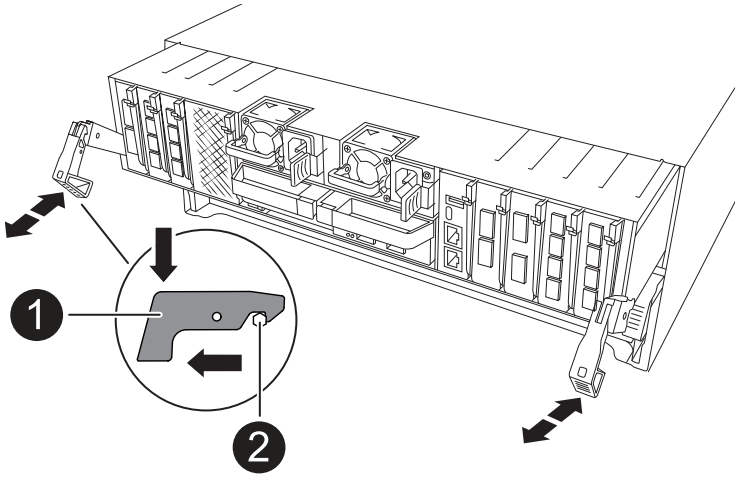

1	Locking latch
2	Locking pin

b. Slide the controller module out of the chassis by the locking latches, and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

c. Repeat these steps for the second controller module.

3. Remove the drives:
   a. Gently remove the bezel from the front of the system.
   b. Press the release button at the top of the drive carrier face below the LEDs.
   c. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      Drives are fragile. Always use two hands to support the drive weight when removing a drive to prevent damage to them.

   d. Keep track of which drive bay each drive was in and set the drive aside on a static-free cart or table.
   e. Repeat this step for the remaining drives in the chassis.

Step 3: Replace the impaired chassis

Remove the impaired chassis and install the replacement chassis.

Steps

1. Remove the impaired chassis:
   a. Remove the screws from the chassis mount points.
   b. Using two people or a lift, slide the impaired chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.

2. Install the replacement chassis:
    a. Using two people or a lift, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
    b. Slide the chassis all the way into the equipment rack or system cabinet.
    c. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.

Step 4: Install the chassis components

After the replacement chassis is installed, you need to install the controller modules, recable them, and then reinstall the drives and PSUs.

Steps

1. Beginning with the bottom controller module, install the controller modules in the replacement chassis:
    a. Align the end of the controller module with the opening in the chassis, and then gently push the controller all the way into the chassis.
    b. Rotate the locking latches upward into the locked position.
    c. If you have not already done so, reinstall the cable management device and recable the controller.

        If you removed the media converters (QSFPs or SFPs), remember to reinstall them.

        Make sure that the cables are connected referencing the cable labels.
2. Reinstall the drives into their corresponding drive bays in the front of the chassis.
3. Install all four of the PSUs:
    a. Using both hands, support and align the edges of the PSU with the opening in the controller module.
    b. Gently push the PSU into the controller module until the locking tab clicks into place.

        The power supplies will only properly engage with the internal connector and lock in place one way.

        > To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

4. Reconnect the PSU power cables to all four of the PSUs.
    a. Secure the power cable to the PSU using the power cable retainer.

        If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis and secure the power cable to the PSU with the thumbscrews.

    The controller modules begin to boot as soon as PSUs are installed and power is restored.

What's next?

After you've replaced the impaired AFF C80 chassis and reinstalled the components into it, you need to complete the chassis replacement.

Complete chassis replacement - AFF C80

Reboot the controllers, verify system health, and return the failed part to NetApp to

complete the final step in the AFF C80 chassis replacement procedure.

Step 1: Boot the controllers and give back the controllers

After the controllers reboot, boot ONTAP and give back the controllers.

Steps

1. Check the console output:
   a. If the controller boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.
   b. If the console displays `waiting for giveback` after the reboot, log into the partner controller and check that the replaced controller is ready for giveback with the `storage failover show` command.
2. Perform the giveback:
   a. Connect the console cable to the partner controller.
   b. Give back the controller with the `storage failover giveback -fromnode local` command.

Step 2: Verify storage system health

After the controller has given back the storage, you should check the overall health with Active IQ Config Advisor.

Steps

1. After the giveback is complete, run Active IQ Config Advisor to verify the health of the storage system.
2. Correct any issues you encounter.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - AFF C80

Get started with replacing the controller in your AFF C80 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

Review the controller replacement requirements

To replace the controller module, you must meet certain requirements.

Shut down the impaired controller

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

3 Replace the controller

Replacing the controller includes removing the impaired controller, moving the FRU components to the replacement controller module, and then installing the replacement controller module in the enclosure.

4 Restore and verify the system configuration

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

5 Recable and give back the controller

Recable the controller and transfer the ownership of storage resources back to the replacement controller.

6 Complete controller replacement

Verify the LIfs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - AFF C80

Before replacing the controller in your AFF C80 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements for replacing the controller module.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- Do not use this procedure for controller upgrades; instead, refer to the Choose your controller hardware upgrade procedure for guidance.
- If your system is in a MetroCluster configuration, you must review Choosing the correct recovery procedure to determine whether you should use this procedure.
- You must replace the failed component with the field-replaceable unit (FRU) you received from NetApp.
- You must replace a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.

- You must always capture the controller's console output to a text log file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

What's next?

After you've reviewed the requirements to replace your AFF C80 controller, you need to shut down the impaired controller.

Shut down the impaired controller - AFF C80

Shut down the controller in your AFF C80 storage system to prevent data loss and ensure system stability when replacing the controller.

Shut down the controller module using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After you've shut down the controller, you need to replace the controller.

Replace the controller - AFF C80

Replace the controller in your AFF C80 system when a hardware failure requires it. This process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting the system.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

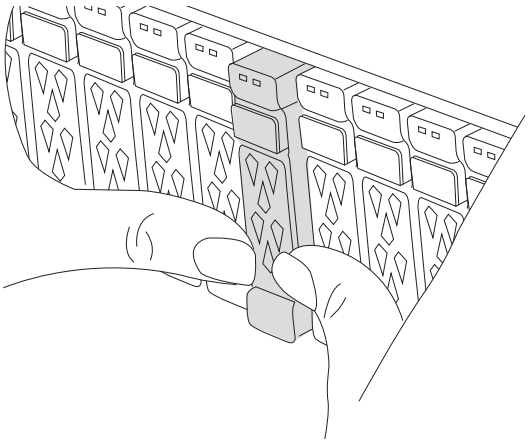

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

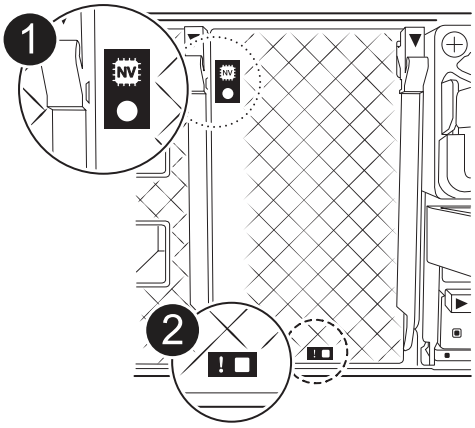

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

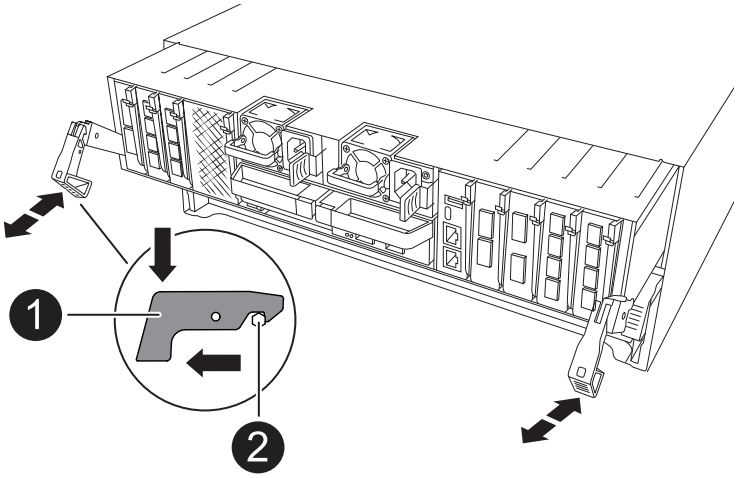

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 2: Move the power supplies

Move the power supplies to the replacement controller.

Steps

1. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

   The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

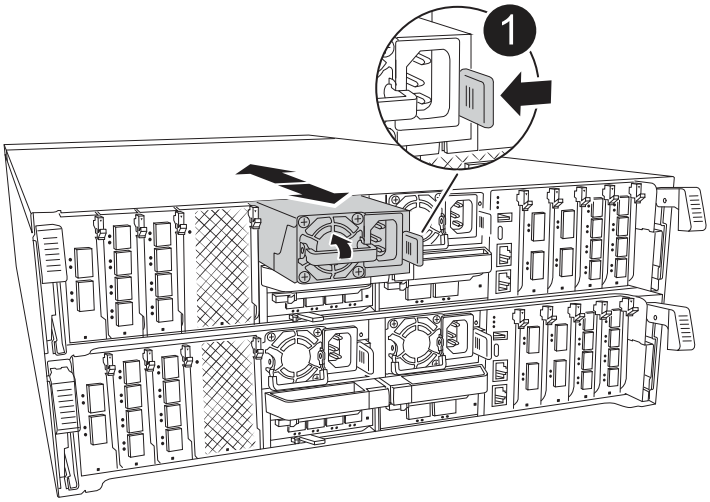

1	Terracotta PSU locking tab
2	Power supply

2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

> To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

Move the fans modules to the replacement controller module.

Steps

1. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

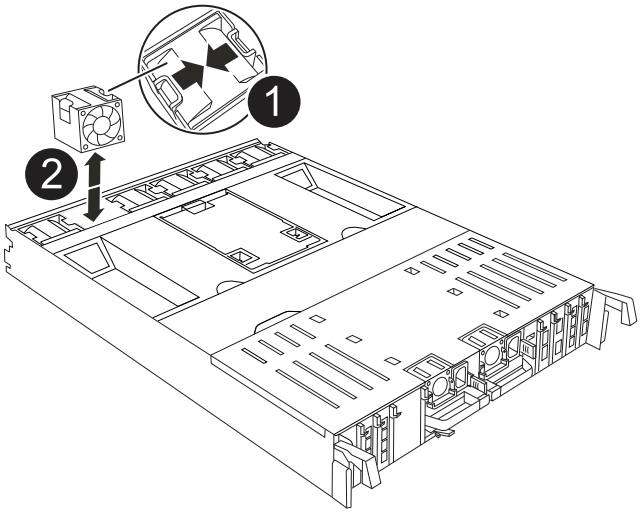

1	Fan locking tabs
2	Fan module

2. Move the fan module to the replacement controller module, and then install the fan module by aligning its edges with the opening in the controller module, and then sliding the fan module into the controller module until the locking latches click into place.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the NV battery

Move the NV battery to the replacement controller module.

Steps

1. Open the air duct cover in the middle of the controller module and locate the NV battery.

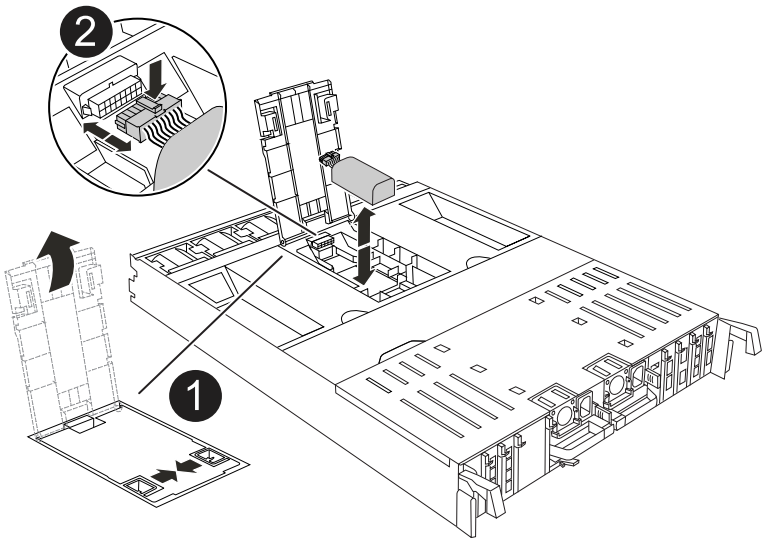

1	NV battery air duct
2	NV battery pack plug

Attention: The NV module LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module.
5. Move the battery pack to the replacement controller module and then install it in the replacement controller module:
   a. Open the NV battery air duct in the replacement controller module.
   b. Plug the battery plug into the socket and make sure that the plug locks into place.
   c. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   d. Close the NV battery air duct.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller module.

Steps

1. Open the controller air duct on the top of the controller.

a. Insert your fingers in the recesses at the far ends of the air duct.

b. Lift the air duct and rotate it upward as far as it will go.

2. Locate the system DIMMs on the motherboard, using the DIMM map on top of the air duct.

   The DIMM locations, by model, are listed in the following table:

Model	DIMM slot location
FAS70	3, 10, 19, 26
FAS90	3, 7, 10, 14, 19, 23, 26, 30

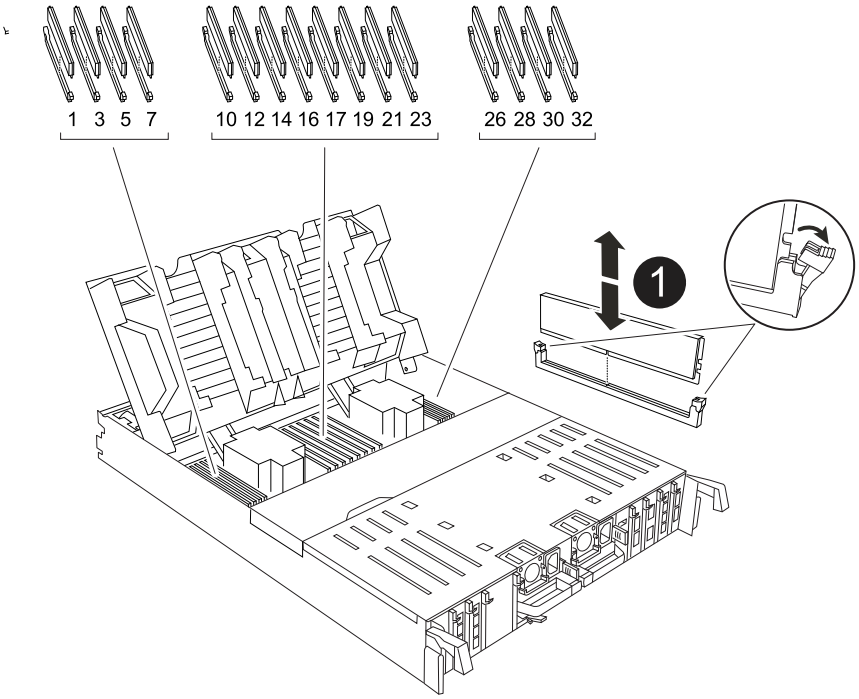

1	System DIMM

3. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.

4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Locate the slot on the replacement controller module where you are installing the DIMM.

6. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Repeat these steps for the remaining DIMMs.
9. Close the controller air duct.

Step 6: Move the I/O modules

Move the I/O modules to the replacement controller module.

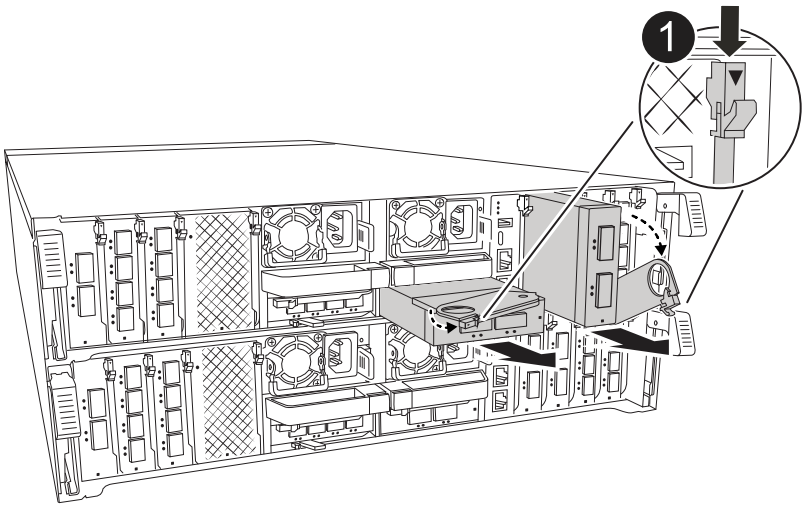

1	I/O module cam lever

Steps

1. Unplug any cabling on the target I/O module.

   Make sure that you label the cables so that you know where they came from.

2. Rotate the cable management arm down by pulling the buttons on the inside of the cable management arm and rotating it down.
3. Remove the I/O modules from the controller module:
   a. Depress the target I/O module cam latch button.
   b. Rotate the cam latch down as far as it will go. For horizontal modules, rotate the cam away from the module as far as it will go.
   c. Remove the module from the controller module by hooking your finger into the cam lever opening and pulling the module out of the controller module.

      Make sure that you keep track of which slot the I/O module was in.

   d. Install the replacement I/O module into the replacement controller module by gently sliding the I/O module into the slot until the I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.
4. Repeat these steps to move the remaining I/O modules, except the modules in slots 6 and 7, to the

replacement controller module.

> To move the I/O modules from slots 6 and 7, you must move the carrier containing these I/O modules from the impaired controller module to the replacement controller module.

5. Move the carrier containing the I/O modules in slots 6 and 7 to the replacement controller module:
   a. Press the button on the right-most handle on the carrier handle.
   ..Slide the carrier out of the impaired controller module insert it into the replacement controller module in the same position it was in the impaired controller module.
   b. Gently push the carrier all the way into the replacement controller module until it locks into place.

Step 7: Move the System Management module

Move the System Management module to the replacement controller module.

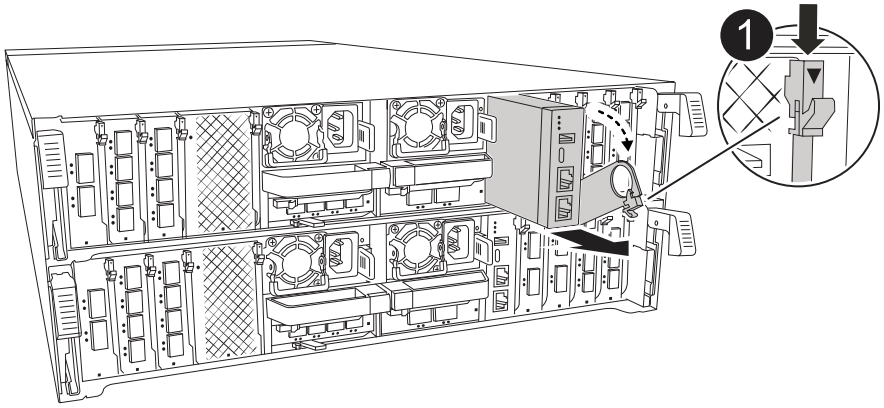

1	System Management module cam latch

Steps

1. Remove the System Management module from the impaired controller module:
   a. Depress the system management cam button.
   b. Rotate the cam lever all the way down.
   c. Loop your finger into the cam lever and pull the module straight out of the system.
2. Install the system management module into the replacement controller module in the same slot that it was in on the impaired controller module:
   a. Align the edges of the System Management module with the system opening and gently push it into the controller module.
   b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

Step 8: Move the NVRAM module

Move the NVRAM module to the replacement controller module.

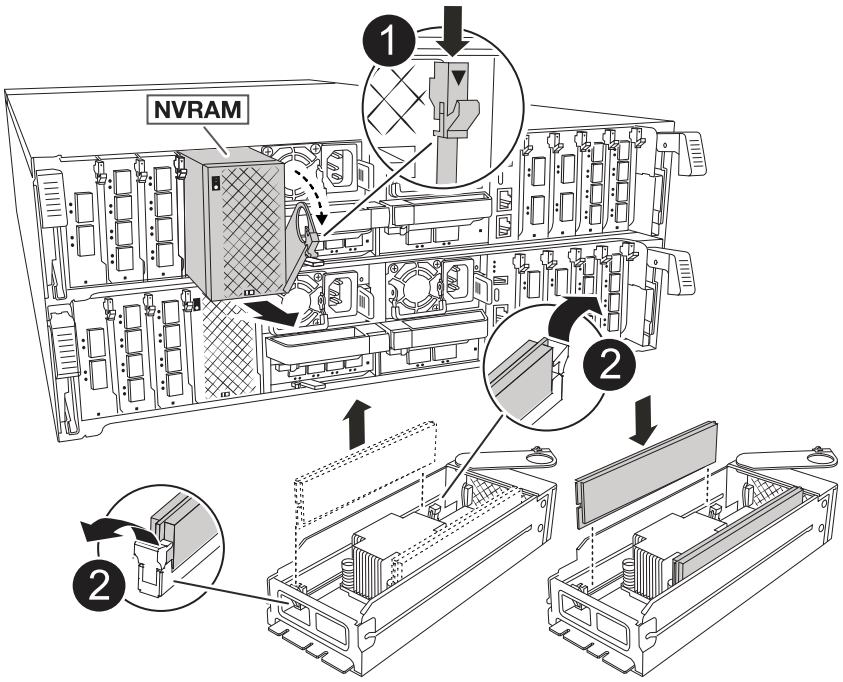

1	Cam locking button
2	DIMM locking tab

Steps

1. Remove the NVRAM module from the impaired controller module:
   a. Depress the cam latch button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.
   c. Remove the NVRAM module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
2. Install the NVRAM module into slot 4/5 in the replacement controller module:
   a. Align the module with the edges of the chassis opening in slot 4/5.
   b. Gently slide the module into the slot all the way, and then push the cam latch all the way up to lock the module in place.

Step 9: Install the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller

module halfway into the system.

> Do not completely insert the controller module in the chassis until instructed to do so.

3. Reinstall the cable management arm, if removed, but do not reconnect any cables to the replacement controller.
4. Plug the console cable into the console port of the replacement controller module and reconnect it to the laptop so that so that it receives console messages when it reboots.
5. Complete the reinstallation of the controller module:

   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

   > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   a. Rotate the locking latches upward into the locked position.
   b. Plug in the power supplies. The controller boots to the LOADER prompt as soon as power is restored.
6. From the LOADER prompt, enter `show date` to display the date and time on the replacement controller. Date and time are in GMT.

   > Time displayed is local time not always GMT and is displayed in 24hr mode.

7. Set the current time in GMT with the `set time hh:mm:ss` command. You can get the current GMT from the partner node the `date -u`command.
8. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

9. Plug the power cords into the power supplies.

   > If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

What's next?

After you've replaced the impaired AFF C80 controller, you need to restore the system configuration.

Restore and verify the system configuration - AFF C80

Verify that the controller's HA configuration is active and functioning correctly in your AFF C80 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

Steps

1. Boot to maintenance mode: `boot_ontap maint`
   a. Enter `y` when you see *Continue with boot?*.

      If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

5. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha`

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed: `ha-config show`

Step 2: Verify disk list

You must verify the adapter list and paths to all your system disks.

Steps

1. Verify that the adapter lists the paths to all disks with the `storage show disk -p`.

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode: `halt`.

What's next?

After you've restored and verified the system configuration for your AFF C80 system, you need to give back the controller.

Give back the controller - AFF C80

Return control of storage resources to the replacement controller so your AFF C80 system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`
   - If you encounter errors, contact NetApp Support.
9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

    c. Verify status of the keys: `security key-manager key query -restored false`

The output should show no results when when properly synchronized.

If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

After you've transferred the ownership of storage resources back to the replacement controller, you need to complete the controller replacement procedure.

Complete controller replacement - AFF C80

To complete the controller replacement for your AFF C80 system, first restore the NetApp Storage Encryption configuration (if necessary). Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, return the failed part to NetApp.

Step 1: Verify LIFs and check cluster health

Before returning the replacement node to service, verify that the logical interfaces are on their home ports, check the cluster health, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any logical interfaces are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF C80

Replace a DIMM in your AFF C80 system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting

ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp..

Steps

1. If you are not already grounded, properly ground yourself.
2. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
3. Locate the DIMMs on your controller module and identify the target DIMM.

   Use the FRU map on the controller airduct to locate the DIMM slot.

4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

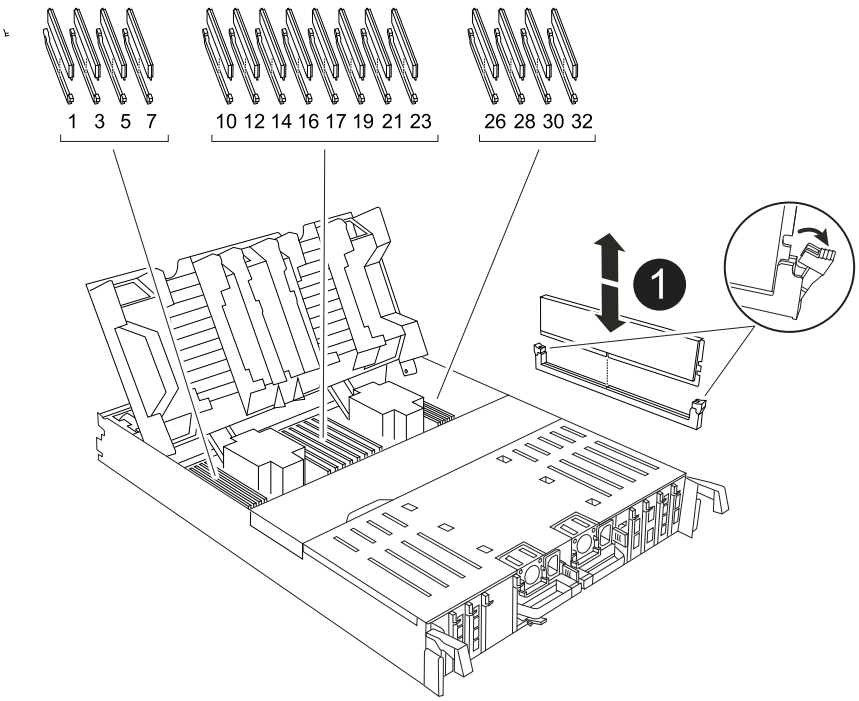

1	DIMM and DIMM ejector tabs

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the controller air duct.

Replace SSD Drive - AFF C80

Replace a drive in your AFF C80 system when a drive fails or requires an upgrade. This process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed SSD drive nondisruptively while I/O is in progress.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed disk drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  It can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the disk depends on how the disk drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

When replacing several disk drives, you must wait 70 seconds between the removal of each failed disk drive and the insertion of the replacement disk drive to allow the storage system to recognize the existence of each new disk.

You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.
8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.

a. Display all unowned drives:

```
storage disk show -container-type unassigned
```

You can enter the command on either controller module.

b. Assign each drive:

```
storage disk assign -disk disk_name -owner node_name
```

You can enter the command on either controller module.

You can use the wildcard character to assign more than one drive at once.

c. Reenable automatic drive assignment if needed:

```
storage disk option modify -node node_name -autoassign on
```

You must reenable automatic drive assignment on both controller modules.

10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Replace a fan module - AFF C80

Replace a fan module in your AFF C80 system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

To replace a fan, remove the failed fan module and replace it with a new fan module.

Steps

1. Identify the fan module that you must replace by checking the console error messages.
2. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

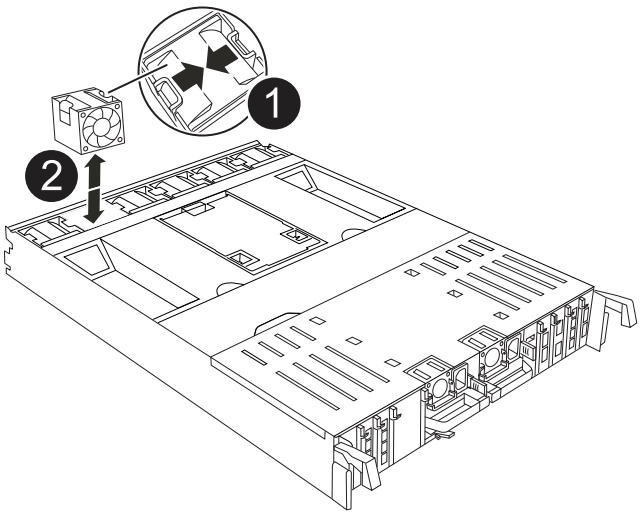

1	Fan locking tabs
2	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module until the locking latches click into place.

Replace NVRAM - AFF C80

Replace the NVRAM in your AFF c80 system when the non-volatile memory becomes faulty or requires an upgrade. The replacement process involves shutting down the impaired controller, replacing the NVRAM module or the NVRAM DIMM, reassigning the disks, and returning the failed part to NetApp.

Replace the NVRAM module or NVRAM DIMMs using the appropriate following option.

Option 1: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 4/5 in the chassis and follow the specific sequence of steps.

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

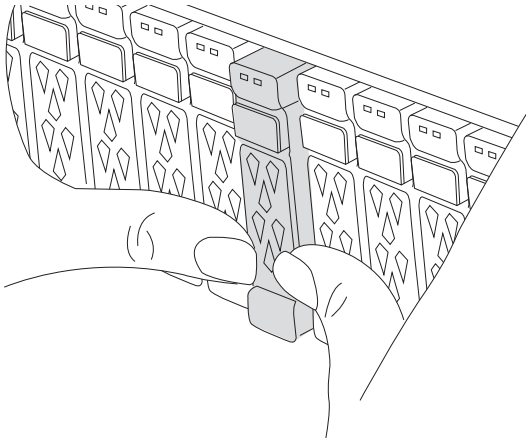

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. If you are not already grounded, properly ground yourself.
4. Unplug the controller's PSUs.

   If your system has DC power, disconnect the power block from the PSUs.

5. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
6. Remove the target NVRAM module from the chassis:
   a. Depress the cam latch button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.
   c. Remove the impaired NVRAM module from the chassis by hooking your finger into the cam lever opening and pulling the module out of the chassis.

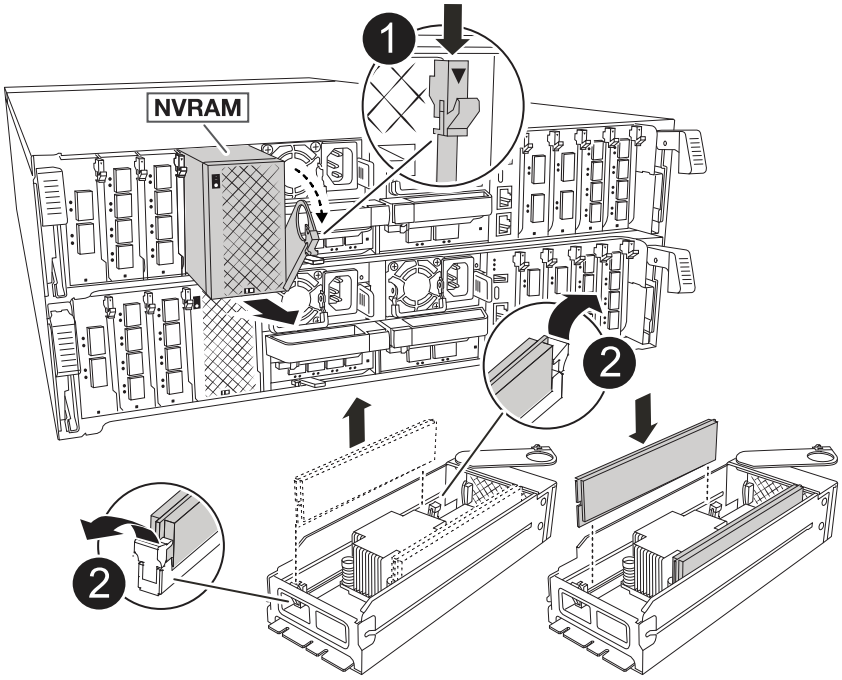

1	Cam locking button
2	DIMM locking tabs

7. Set the NVRAM module on a stable surface.
8. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
9. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 4/5.
   b. Gently slide the module into the slot all the way, and then push the cam latch all the way up to lock the module in place.
10. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

    If you have DC power supplies, reconnect the power block to the power supplies.

11. Rotate the cable management tray up to the closed position.
12. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*.
13. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`.
14. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Option 2: Replace the NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, first must remove the NVRAM module and then replace the target DIMM.

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

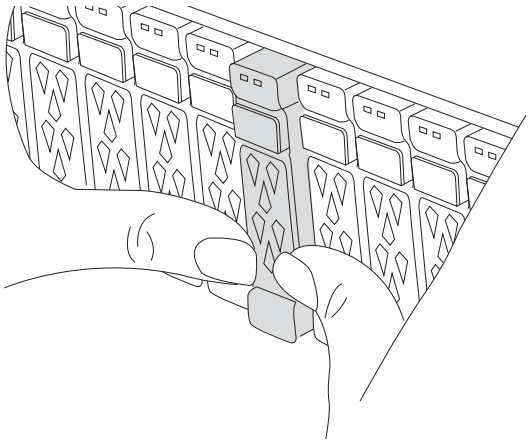

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller's PSUs.

   If your system has DC power, disconnect the power block from the PSUs.

4. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
5. Remove the target NVRAM module from the chassis:
   a. Depress the cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.
   c. Remove the NVRAM module from the chassis by hooking your finger into the cam lever opening and pulling the module out of the chassis.

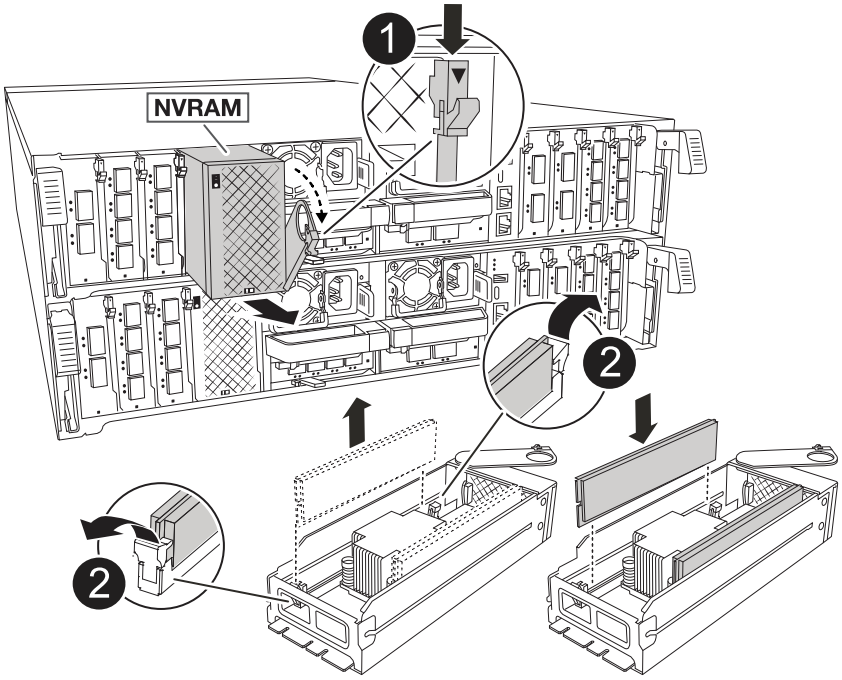

1	Cam locking button
2	DIMM locking tabs

6. Set the NVRAM module on a stable surface.
7. Locate the DIMM to be replaced inside the NVRAM module.

> Consult the FRU map label on the side of the NVRAM module to determine the locations of DIMM slots 1 and 2.

8. Remove the DIMM by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
9. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
10. Install the NVRAM module into the chassis:
    a. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.
11. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

> If you have DC power supplies, reconnect the power block to the power supplies.

12. Rotate the cable management tray up to the closed position.
13. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*.
14. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`.

15. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 3: Reassign disks

You must confirm the system ID change when you boot the controller and then verify that the change was implemented.

> Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt:

   `halt`

2. From the LOADER prompt on the controller, boot the controller and enter *y* when prompted to override the system ID due to a system ID mismatch.
3. Wait until the Waiting for giveback message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify that the new partner system ID has been automatically assigned:

   `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node 2 has undergone replacement and has a new system ID of 151759706.

```
node1:> storage failover show
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage:

      `storage failover giveback -ofnode replacement_node_name`

      The controller takes back its storage and completes booting.

If you are prompted to override the system ID due to a system ID mismatch, you should enter *y*.

If the giveback is vetoed, you can consider overriding the vetoes.

For more information, see manual giveback commands to override the veto.

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: *storage failover show*

The output from the `storage failover show` command should not include the System ID changed on partner message.

5. Verify that the disks were assigned correctly:

`storage disk show -ownership`

The disks belonging to the controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 151759706:

```
node1:> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        151759706  151759706  -
151759706 Pool0
1.0.1  aggr0_1  node1 node1           151759706  151759706  -
151759706 Pool0
.
.
.
```

6. If the system is in a MetroCluster configuration, monitor the status of the controller: *metrocluster node show*

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

7. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The controller is the current owner of the disks on the disaster site.

See Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration for more information.

8. If your system is in a MetroCluster configuration, verify that each controller is configured: *metrocluster node show - fields configuration-state*

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

9. Verify that the expected volumes are present for each controller:

   `vol show -node node-name`

10. If storage encryption is enabled, you must restore functionality.
11. Return the impaired controller to normal operation by giving back its storage:

    `storage failover giveback -ofnode impaired_node_name.`

12. If automatic giveback was disabled, reenable it:

    `storage failover modify -node local -auto-giveback true.`

13. If AutoSupport is enabled, restore/unsuppress automatic case creation:

    `system node autosupport invoke -node * -type all -message MAINT=END.`

Replace the NV battery - AFF C80

Replace the NV battery in your AFF C80 system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

Remove the failed NV battery from the controller module and install the replacement NV battery.

Steps

1. Open the air duct cover and locate the NV battery.

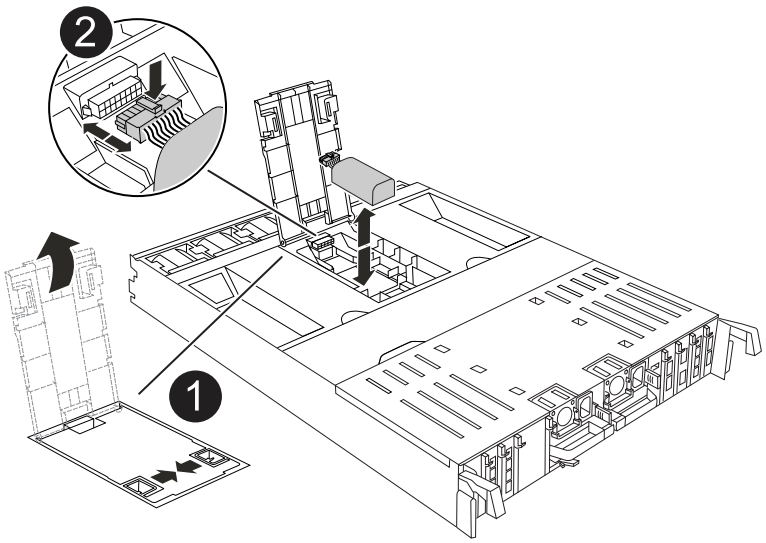

1	NV battery air duct cover
2	NV battery plug

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module, and then set it aside.
5. Remove the replacement battery from its package.
6. Install the replacement battery pack into the controller:
   a. Plug the battery plug into the riser socket and make sure that the plug locks into place.
   b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
7. Close the NV air duct cover.

   Make sure that the plug locks into the socket.

I/O module

Overview of add and replace an I/O module - AFF C80

The AFF C80 system offers flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your AFF C80 storage system with the same type of I/O module, or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

- Add an I/O module

  Adding additional modules can improve redundancy, helping to ensure that the system remains operational

even if one module fails.

- Replace an I/O module

  Replacing a failing I/O module can restore the system to its optimal operating state.

I/O slot numbering

The I/O slots on AFF C80 controllers are numbered 1 through 11, as shown in the following illustration.

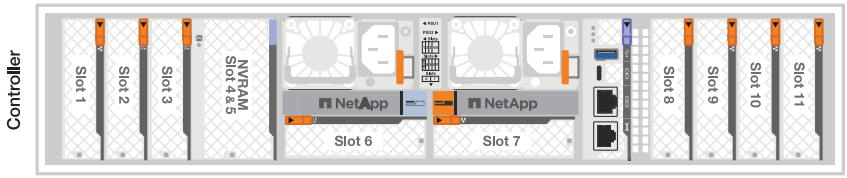

Add an I/O module - AFF C80

Add an I/O module to your AFF C80 system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your AFF C80 storage system when there are empty slots available or when all slots are fully populated.

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- Make sure that all other components are functioning properly.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the target slot blanking module from the carrier:
   a. Depress the cam latch on the blanking module in the target slot.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
4. Install the I/O module:
   a. Align the I/O module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
5. Cable the I/O module to the designated device.

   Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

6. Rotate the cable management tray up to the closed position.
7. From the LOADER prompt, reboot the node:

   ```
   bye
   ```

   This reinitializes the I/O module and other components and reboots the node.

8. Give back the controller from the partner controller:

   ```
   storage failover giveback -ofnode target_node_name
   ```

9. Repeat these steps for controller B.
10. From the healthy node, restore automatic giveback if you disabled it:

    ```
    storage failover modify -node local -auto-giveback true
    ```

11. If AutoSupport is enabled, restore automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the target I/O module from the chassis:
   a. Depress the cam latch button.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

      Make sure that you keep track of which slot the I/O module was in.
5. Install the I/O module into the target slot in the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
6. Cable the I/O module to the designated device.
7. Repeat the remove and install steps to replace additional modules for the controller.
8. Rotate the cable management tray up to the closed position.
9. Reboot the controller from the LOADER prompt:_bye_

   This reinitializes the PCIe cards and other components and reboots the node.

   If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement
10. Give back the controller from the partner controller:

    ```
    storage failover giveback -ofnode target_node_name
    ```
11. Enable automatic giveback if it was disabled:

```
storage failover modify -node local -auto-giveback true
```

12. Do one of the following:
    - If you removed a NIC I/O module and installed a new NIC I/O module, use the following network command for each port:

      ```
      storage port modify -node *<node name> -port *<port name> -mode network
      ```

    - If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
13. Repeat these steps for controller B.

Replace an I/O module - AFF C80

Replace an I/O module in your AFF C80 system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

To replace an I/O module, locate it within the controller module and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.

   Make sure to label the cables so that you know where they came from.
3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the I/O module from the controller module:

   This following illustration shows removing a horizontal and vertical I/O module. Typically, you will only remove one I/O module.

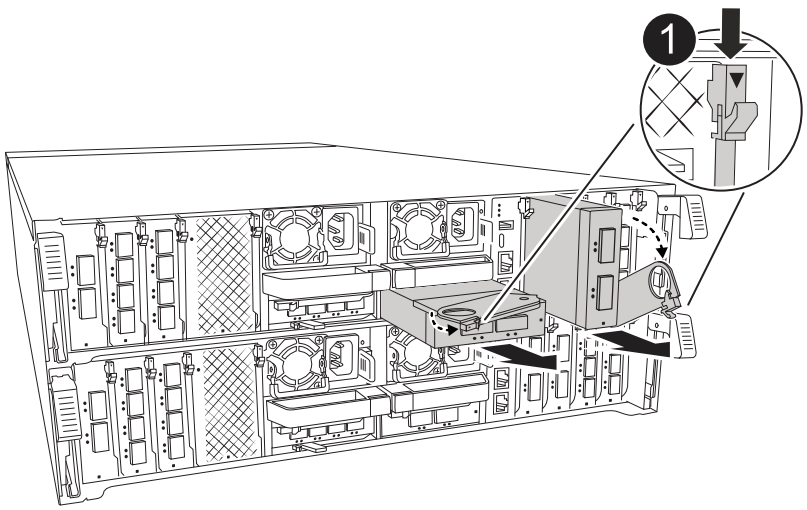

1	Cam locking button

a. Depress the cam latch button.
b. Rotate the cam latch do away from the module as far as it will go.
c. Remove the module from the controller module by hooking your finger into the cam lever opening and pulling the module out of the controller module.

   Make sure that you keep track of which slot the I/O module was in.

5. Set the I/O module aside.
6. Install the replacement I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently slide the module into the slot all the way into the controller module, and then rotate the cam latch all the way up to lock the module in place.
7. Cable the I/O module.
8. Repeat the remove and install steps to replace additional modules for the controller.
9. Rotate the cable management tray into the locked position.

Replace a power supply - AFF C80

Replace an AC or DC power supply unit (PSU) in your AFF C80 system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cable, replacing the faulty PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

About this task

This procedure is written for replacing one PSU at a time.

> Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

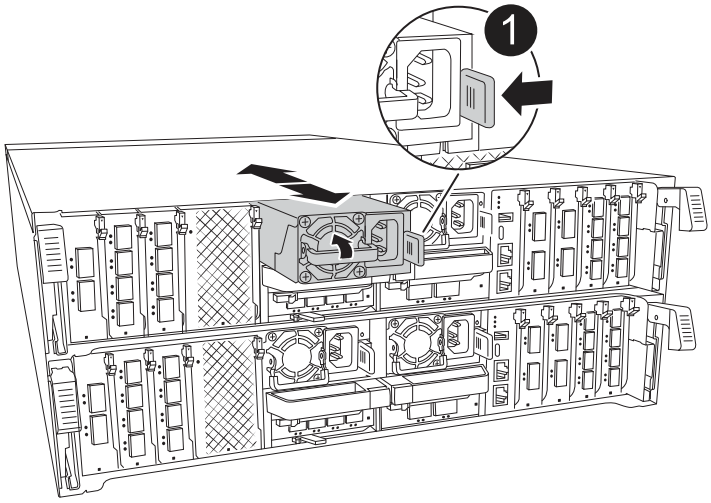

1	Terracotta PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.
6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU.

b. Secure the power cable to the PSU using the power cable retainer.

Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

> The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

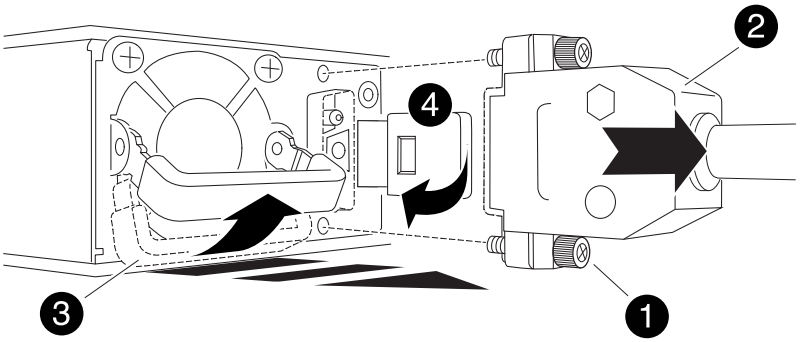

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle
4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.

b. Gently push the PSU into the controller module until the locking tab clicks into place.

The power supplies will only properly engage with the internal connector and lock in place one way.

> To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

   Once power is restored to the PSU, the status LED should be green.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF C80

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your AFF C80 system to ensure that services and applications relying on accurate time synchronization remain operational.

Remove the failed RTC battery and install the replacement RTC battery.

Steps

1. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the RTC battery under the air duct.

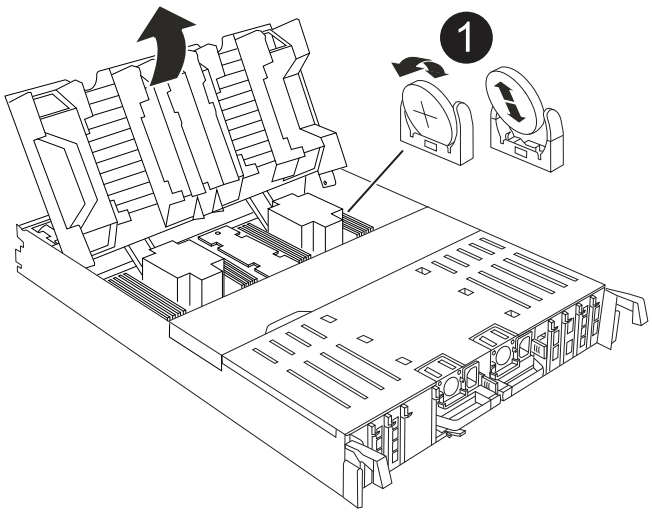

1	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Replace system management module - AFF C80

Replace the System Management module in your AFF C80 system when it becomes defective or its firmware is corrupted. The replacement process involves shutting down the controller, replacing the failed System Management module, rebooting the controller, updating the license keys, and returning the failed part to NetApp.

Replace the impaired system management module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

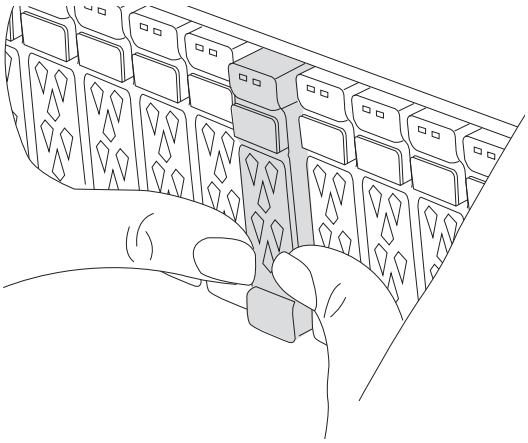

2. Make sure NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged. If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

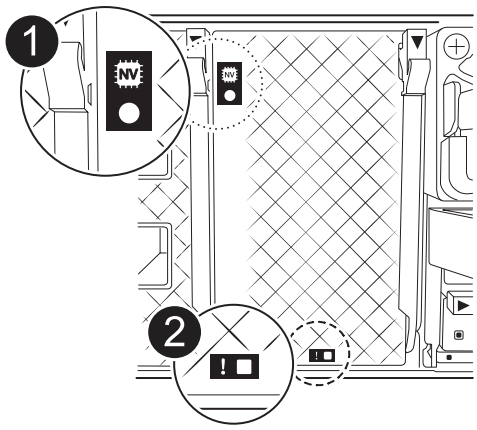

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
4. Unplug the controller's PSUs.

> If your system has DC power, disconnect the power block from the PSUs.

5. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
6. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.

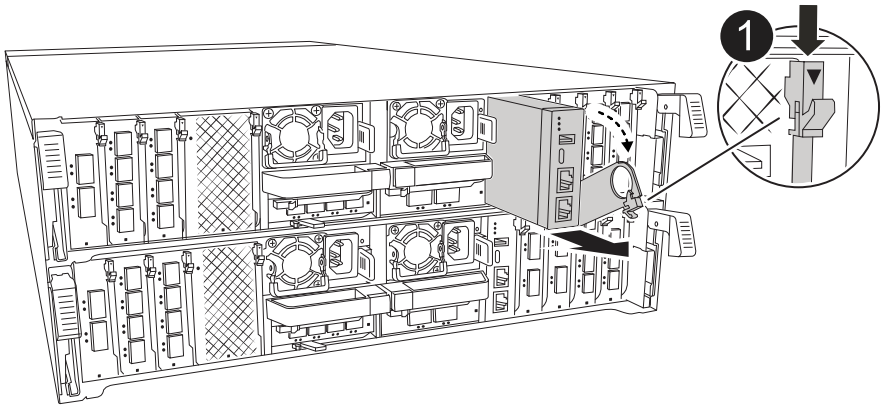

1	System Management module cam latch

7. Remove the System Management module:

a. Depress the system management cam button.
The cam lever moves away from the chassis.

b. Rotate the cam lever all the way down.

c. Loop your finger into the cam lever and pull the module straight out of the system.

d. Place the System Management module on an anti-static mat, so that the boot media is accessible.

8. Move the boot media to the replacement System Management module:

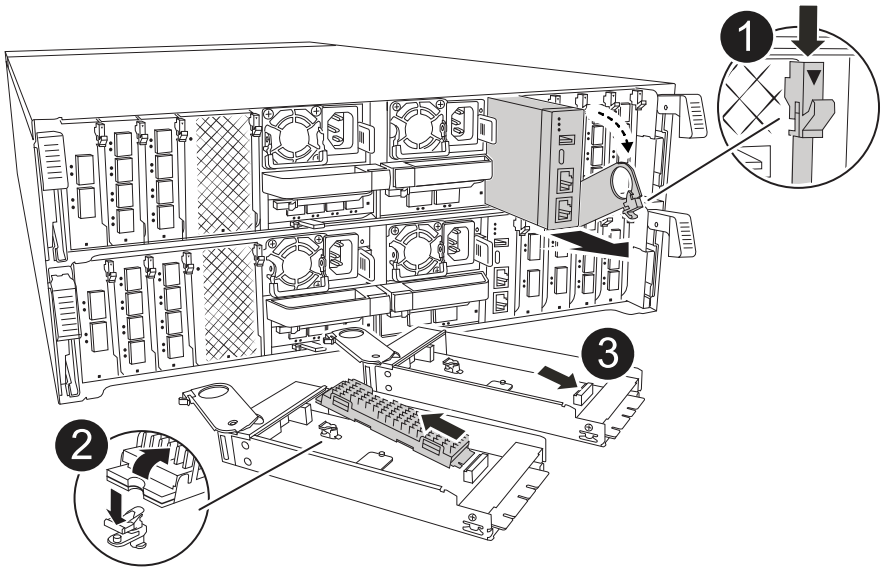

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue locking button.
The boot media rotates slightly upward.

b. Rotate the boot media up, slide it out of the socket.

c. Install the boot media in the replacement System Management module:

i. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

ii. Rotate the boot media down toward until it engages the locking button. Depress the blue locking if necessary.

9. Install the system management module:

a. Align the edges of the replacement System Management module with the system opening and gently push it into the controller module.

b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

10. Recable the System Management module.

11. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

    If you have DC power supplies, reconnect the power block to the power supplies.

12. Rotate the cable management tray up to the closed position.

AFF C250 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick steps - AFF C250

This section gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Access the *Installation and Setup Instructions* PDF poster:

AFF C250 Installation and Setup Instructions

Video steps - AFF C250

The following video shows how to install and cable your new system.

Animation - Installation and Setup of an AFF C250

If you have a MetroCluster configuration, use the MetroCluster installation content.

MetroCluster Documentation

Detailed steps - AFF C250

This procedure gives detailed step-by-step instructions for installing an AFF C250 storage system.

If you have a MetroCluster configuration, use the MetroCluster Documentation.

Step 1: Prepare for installation

To install your AFF C250 system, you need to create an account and register the system. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

- Make sure you have access to the NetApp Hardware Universe (HWU) for information about site requirements.
- Customers with specific power requirements must check HWU for configuration options.
- Make sure you have access to the Release Notes for your version of ONTAP for more information about this system.
- You need to provide the following at your site:
  - Rack space for the storage system
  - Phillips #2 screwdriver
  - Additional networking cables to connect your system to your network switch and laptop or console with a Web browser.

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register (NetApp Product Registration) your system.
4. Download and install NetApp Downloads: Config Advisor on your laptop.
5. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
25 GbE cable	X66240A-05 (112-00595), 0.5m; X66240-2 (112-00573), 2m		Cluster interconnect network
	X66240A-2 (112-00598), 2m; X66240A-5 (112-00600), 5m		Data
100 GbE cable	X66211-2 (112-00574), 2m; X66211-5 (112-00576), 5m		Storage

Type of cable…	Part number and length	Connector type	For…
RJ-45 (order dependent)	Not applicable		Management network (BMC and wrench port) and Ethernet data (e0a and e0b)
Fibre Channel	X66250-2 (112-00342) 2m; X66250-5 (112-00344) 5m; X66250-15 (112-00346) 15m; X66250-30 (112-00347) 30m		
Micro-USB console cable	Not applicable		Console connection during software setup
Power cables	Not applicable		Powering up the system

6. Review the ONTAP Configuration Guide and collect the required information listed in that guide.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

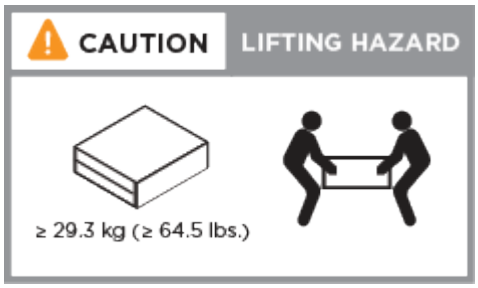

3. Identify and manage cables because this system does not have a cable management device.
4. Place the bezel on the front of the system.

Step 3: Cable controllers to cluster

Cable the controllers to a cluster by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Two-node switchless cluster

The management, Fibre Channel, and data or host network ports on the controller modules are connected to switches. The cluster interconnect ports are cabled on both controller modules.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

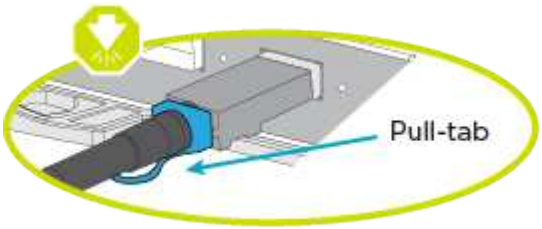

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the tabulated steps to complete the cabling between the controllers and the switches. Perform the steps on each controller.

Animation - Cable a two-node switchless cluster

Steps

1. Cable the cluster interconnect ports e0c to e0c and e0d to e0d with the 25GbE cluster interconnect cables.


:

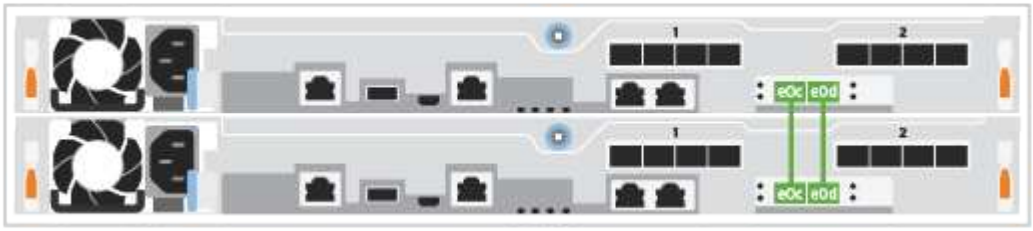

2. Cable the wrench ports to the management network switches with the RJ45 cables.

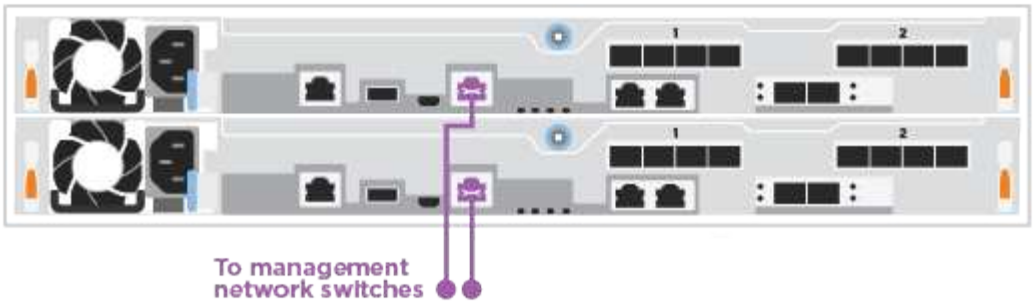

DO NOT plug in the power cords at this point.

Option 2: Switched cluster

All ports on the controllers are connected to switches; cluster interconnect, management, Fibre Channel, and data or host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

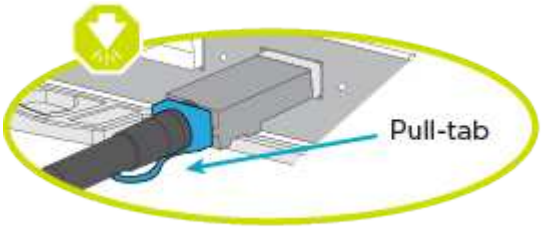

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the tabulated steps to complete the cabling between the controllers and the switches. Perform the steps on each controller.

Animation - Cable a switched cluster

Steps

1. Cable the cluster interconnect ports e0c and e0d to the 25 GbE cluster interconnect switches.

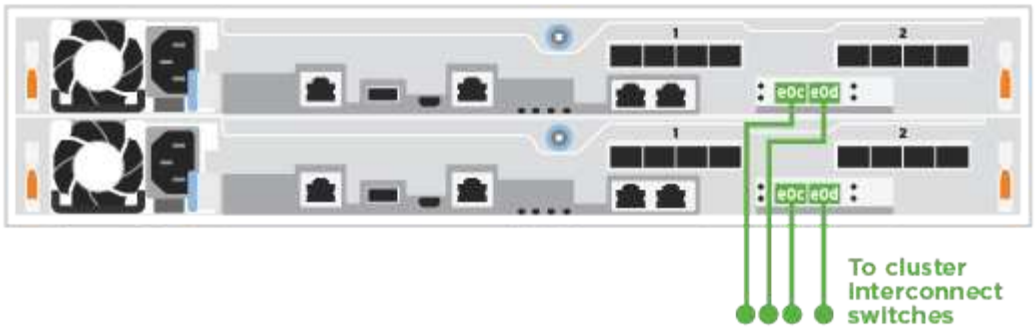

2. Cable the wrench ports to the management network switches with the RJ45 cables.

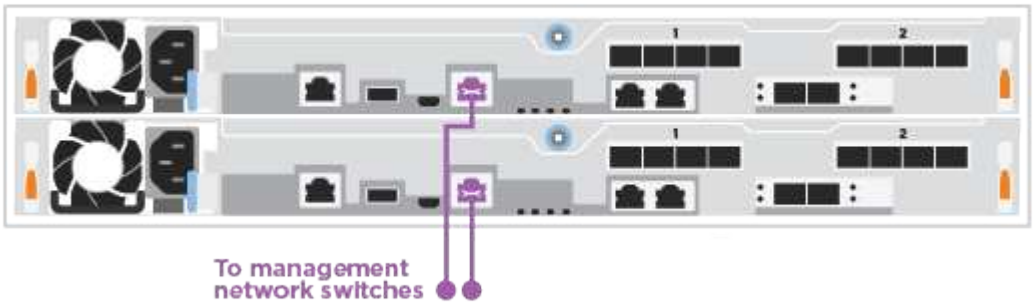

DO NOT plug in the power cords at this point.

Step 4: Cable to host network or storage (Optional)

You have configuration-dependent optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cabling to a host network and storage.

NetApp Hardware Universe slot priority for host network cards (Fibre Channel or 25GbE) is slot 2. However, if you have both cards, the Fibre Channel card goes in slot 2 and the 25GbE card goes in slot 1 (as shown in the options below). If you have an external shelf, the storage card goes in slot 1, the only supported slot for shelves.

Option 1: Cable to Fibre Channel host network

Fibre Channel ports on the controllers are connected to Fibre Channel host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

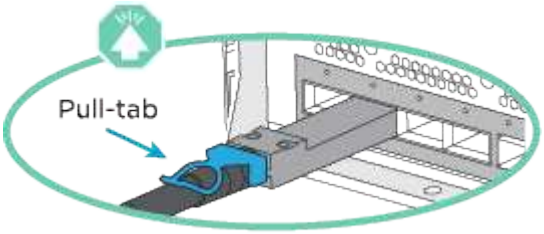

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Perform the step on each controller module.

Steps

1. Cable ports 2a through 2d to the FC host switches.

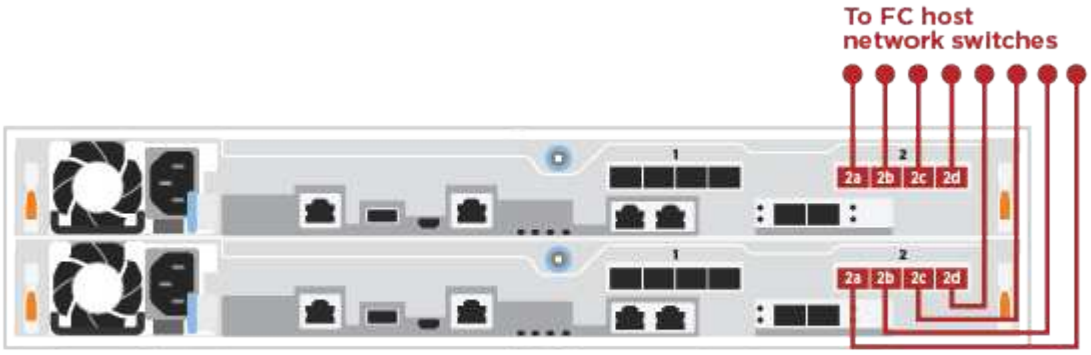

Option 2: Cable to a 25GbE data or host network

25GbE ports on the controllers are connected to 25GbE data or host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

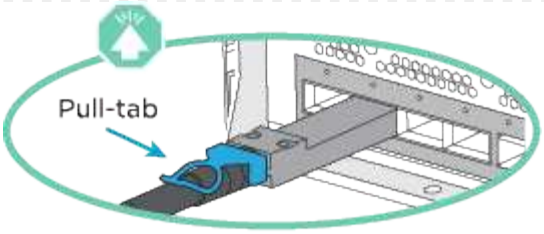

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Perform the step on each controller module.

Steps

1. Cable ports e4a through e4d to the 10GbE host network switches.

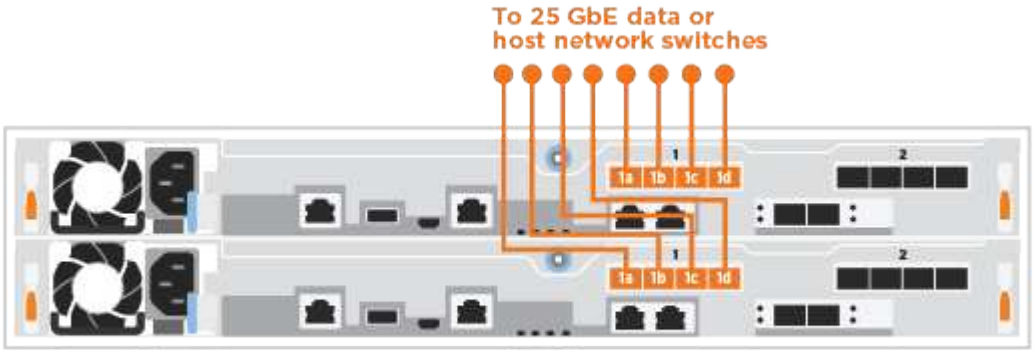

Option 3: Cable controllers to single drive shelf

Cable each controller to the NSM modules on the NS224 drive shelf.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

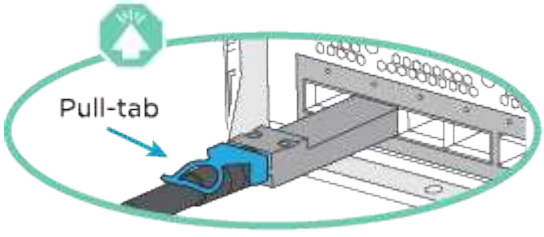

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the tabulated steps to complete the cabling between the controllers and the single shelf. Perform the steps on each controller module.

Animation - Cable the controllers to a single NS224

Steps

1. Cable controller A to the shelf.

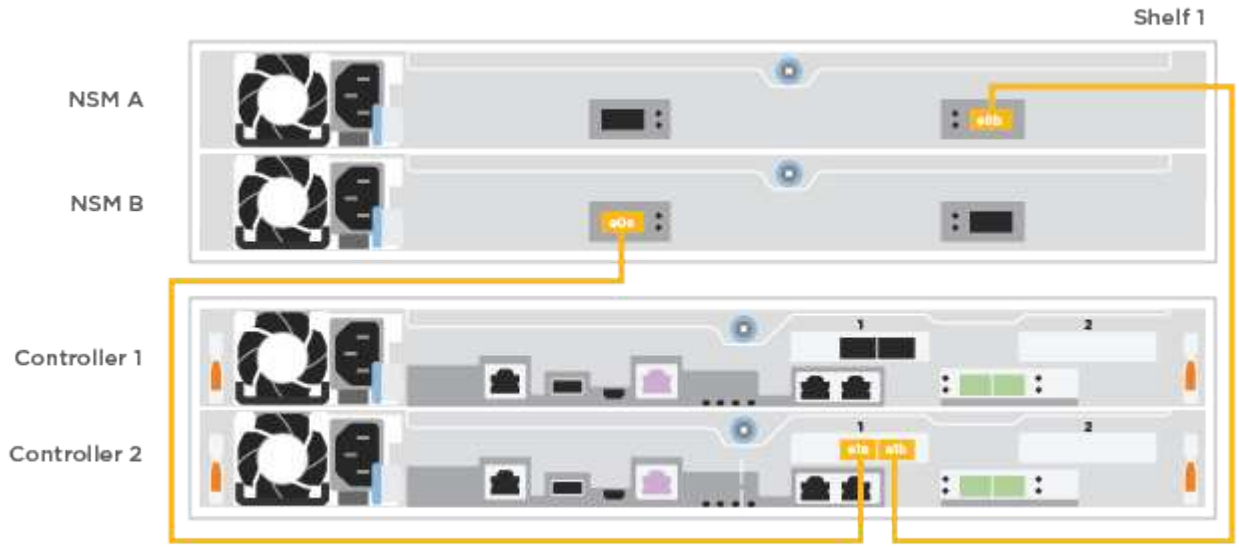

2. Cable controller B to the shelf.

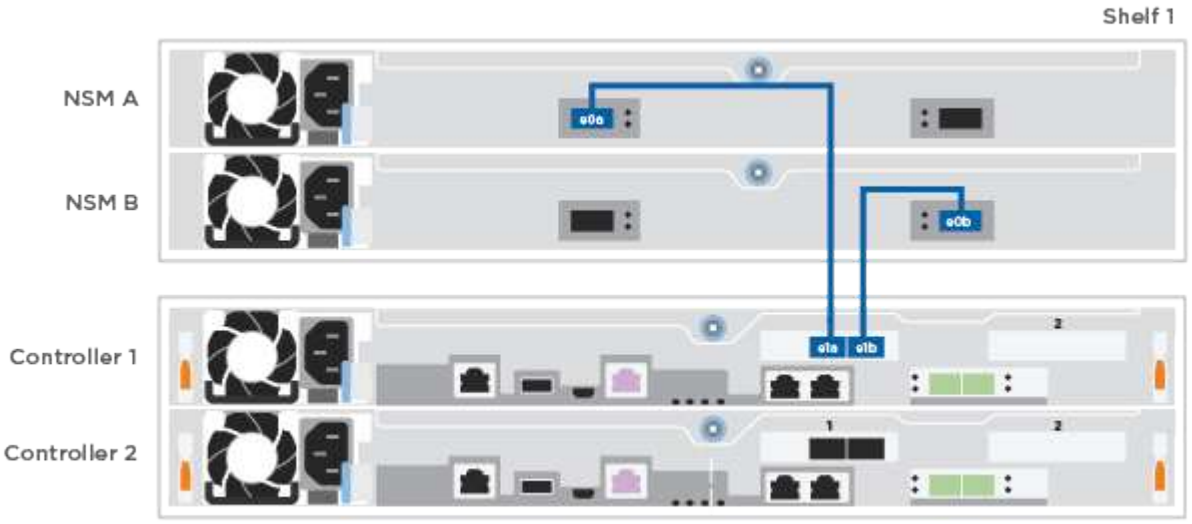

Step 5: Complete system setup

Complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: If network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Use the following animation to power on and set shelf IDs for one or more drive shelves:

   For NS224 drive shelves, shelf IDs are pre-set to 00 and 01. If you want to change the shelf IDs, use the straightened end of a paperclip, or narrow tipped ball point pen to access the shelf ID button behind the faceplate.

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Connect your laptop to the Management switch:

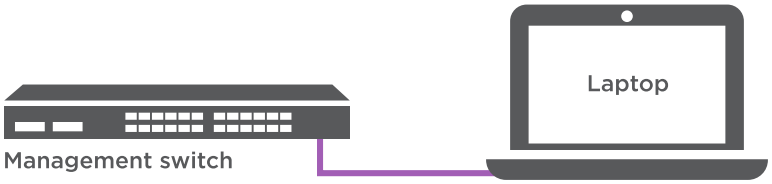

5. Select an ONTAP icon listed to discover:

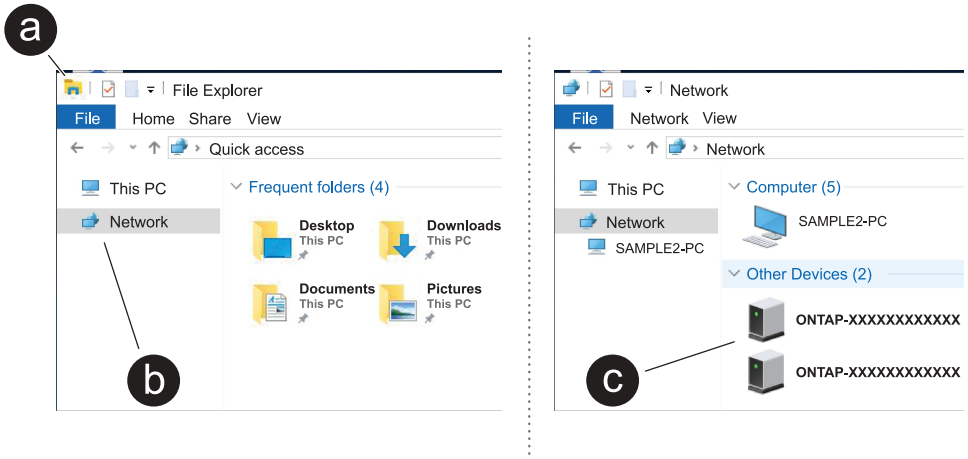

   a. Open File Explorer.

   b. Click **Network** in the left pane and right-click and select **refresh**.

   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

   XXXXX is the system serial number for the target node.

System Manager opens.

6. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the laptop or console to the management switch.

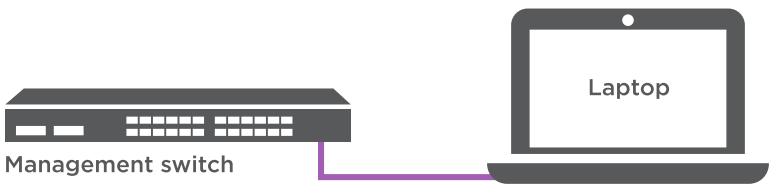

   c. Assign a TCP/IP address to the laptop or console, using one that is on the management switch.
2. Use the following animation to power on and set shelf IDs for one or more drive shelves:

   For NS224 drive shelves, shelf IDs are pre-set to 00 and 01. If you want to change the shelf IDs, use the straightened end of a paperclip, or narrow tipped ball point pen to access the shelf ID button behind the faceplate.

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

The system begins to boot. Initial booting may take up to eight minutes.

4. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.
   The format for the address is https://x.x.x.x.
   b. Configure the system using the data you collected in the ONTAP Configuration Guide.
6. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.
   b. Register your system.
   c. Download Active IQ Config Advisor.
7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain AFF C250 hardware

For the AFF C250 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

Fan

The fan cools the controller.

Mezzanine card

A Mezzanine card is a printed circuit board that plugs directly into another plug-in card.

NVMEM battery

A battery is included with the controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF C250

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.
- You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

About this task

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but

the process requires two reboots.

- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
  - The *impaired* node is the controller on which you are performing maintenance.
  - The *healthy* node is the HA partner of the impaired controller.

Check encryption key support and status - AFF C250

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.

ONTAP version	Run this command
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller - AFF C250

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller displays…	Then…
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Option 2: Systems in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *impaired_node_name* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - AFF C250

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller module, you must first remove the controller module from the system, and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Unplug the I/O cables from the controller module.
5. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

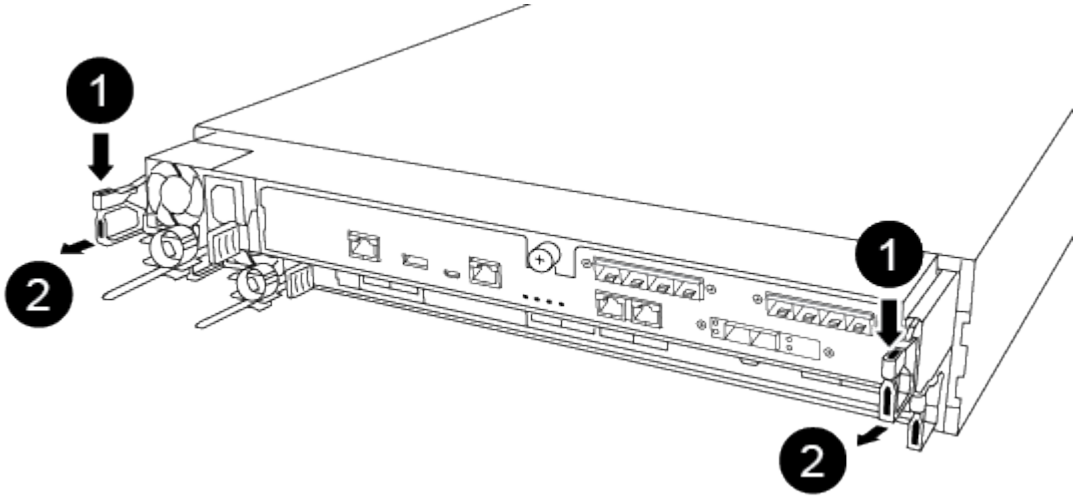

1	Lever
2	Latching mechanism

6. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
7. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

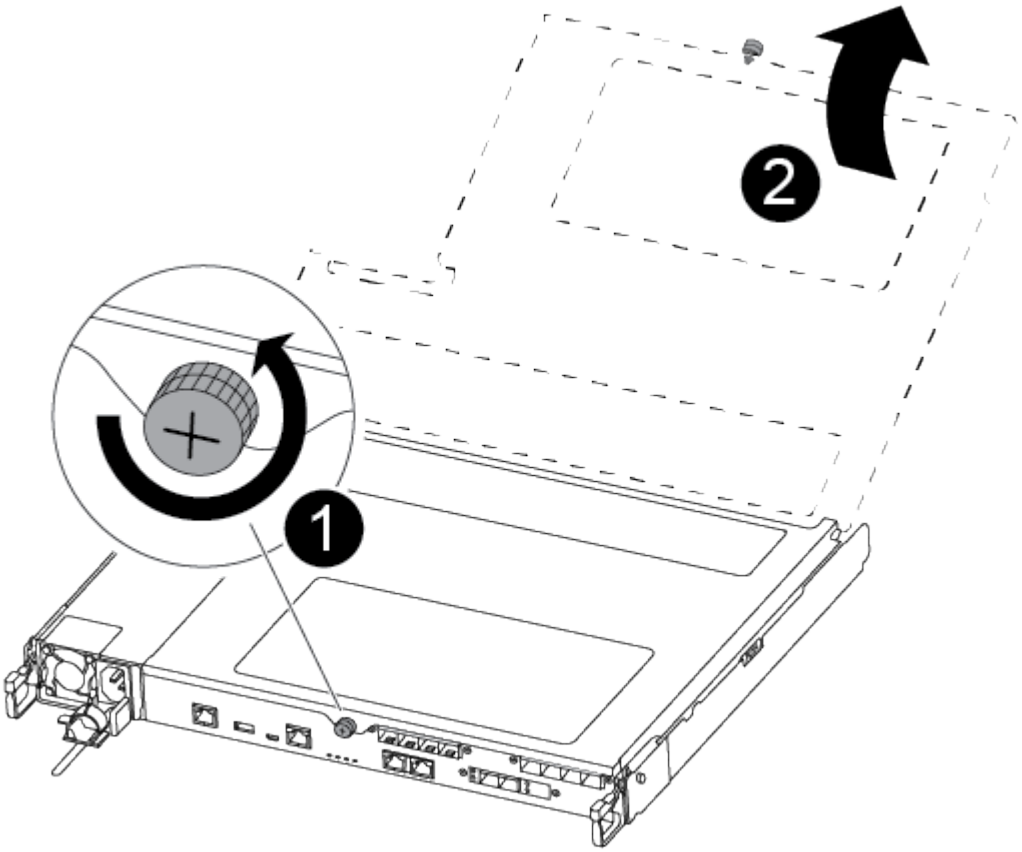

1	Thumbscrew
2	Controller module cover.

8. Lift out the air duct cover.

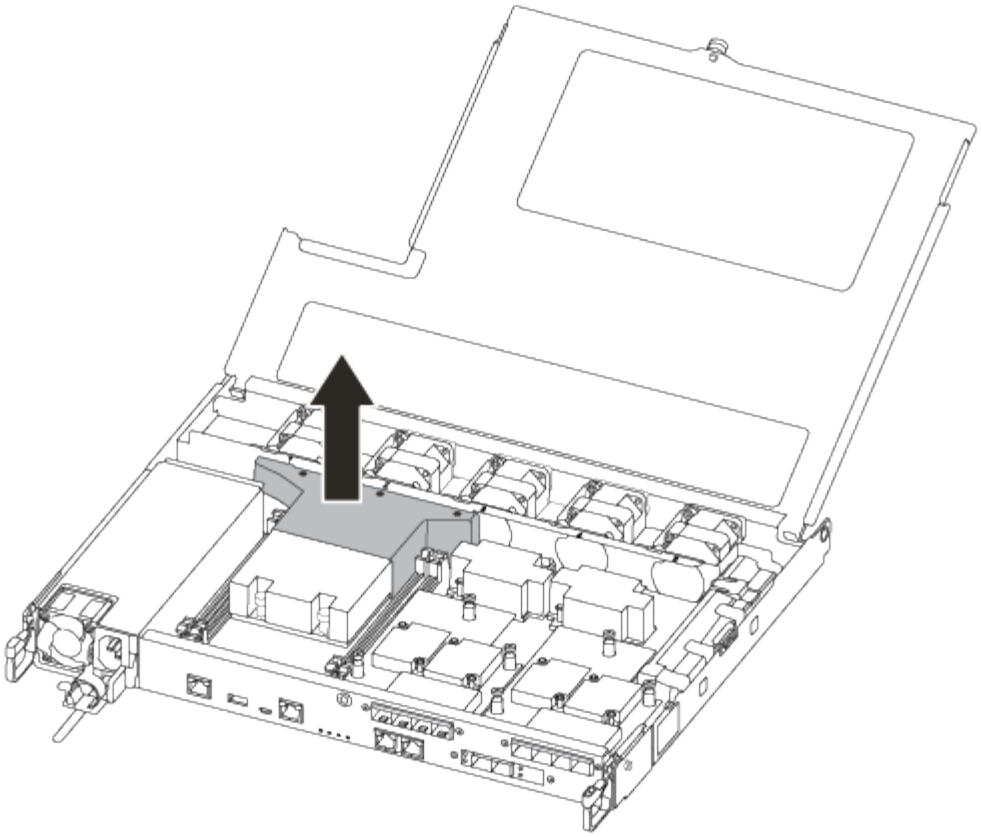

Step 2: Replace the boot media

You locate the failed boot media in the controller module by removing the air duct on the controller module before you can replace the boot media.

You need a #1 magnetic Phillips head screwdriver to remove the screw that holds the boot media in place. Due to the space constraints within the controller module, you should also have a magnet to transfer the screw on to so that you do not lose it.

You can use the following video or the tabulated steps to replace the boot media:

Animation - Replace the boot media

1. Locate and replace the impaired boot media from the controller module.

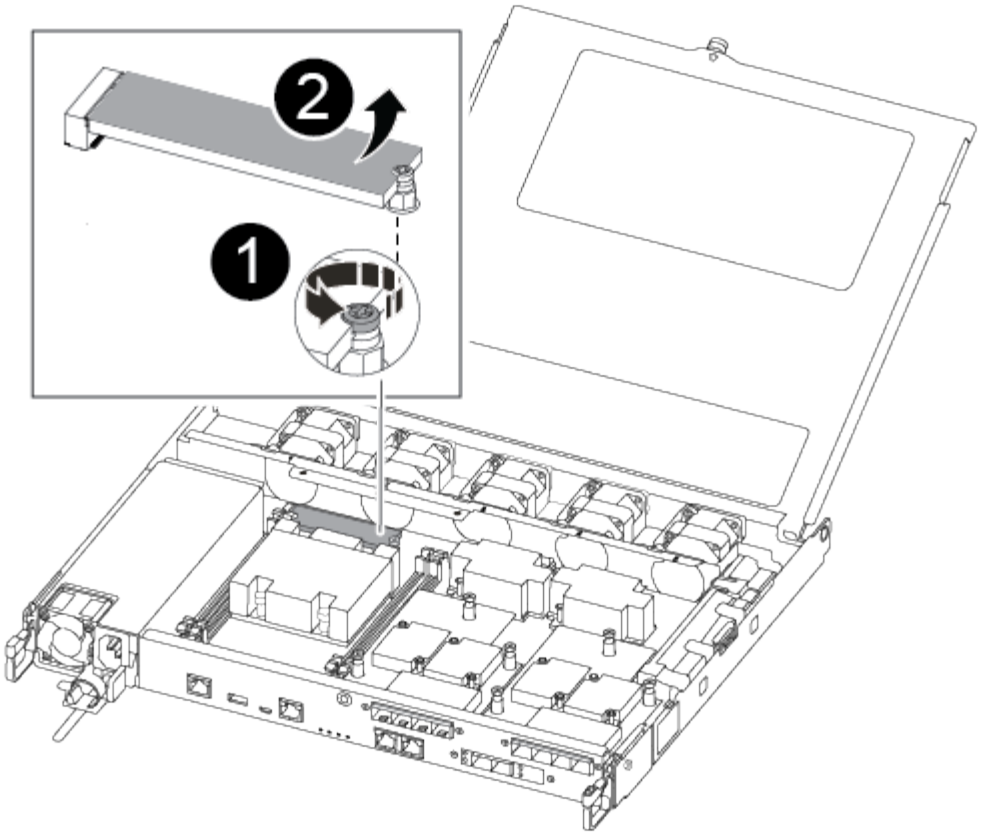

1	Remove the screw securing the boot media to the motherboard in the controller module.
2	Lift the boot media out of the controller module.

2. Using the #1 magnetic screwdriver, remove the screw from the impaired boot media, and set it aside safely on the magnet.
3. Gently lift the impaired boot media directly out of the socket and set it aside.
4. Remove the replacement boot media from the antistatic shipping bag and align it into place on the controller module.
5. Using the #1 magnetic screwdriver, insert and tighten the screw on the boot media.

> Do not apply force when tightening the screw on the boot media; you might crack it.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed is without a boot image so you need to transfer a boot image using a USB flash drive.

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download

button.

  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
2. Download the service image to your work space on your laptop.
3. Unzip the service image.

   If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

   There are two folders in the unzipped service image file:

   - boot
   - efi

4. Copy the efi folder to the top directory on the USB flash drive.

   If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

   The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

5. Remove the USB flash drive from your laptop.
6. If you have not already done so, install the air duct.

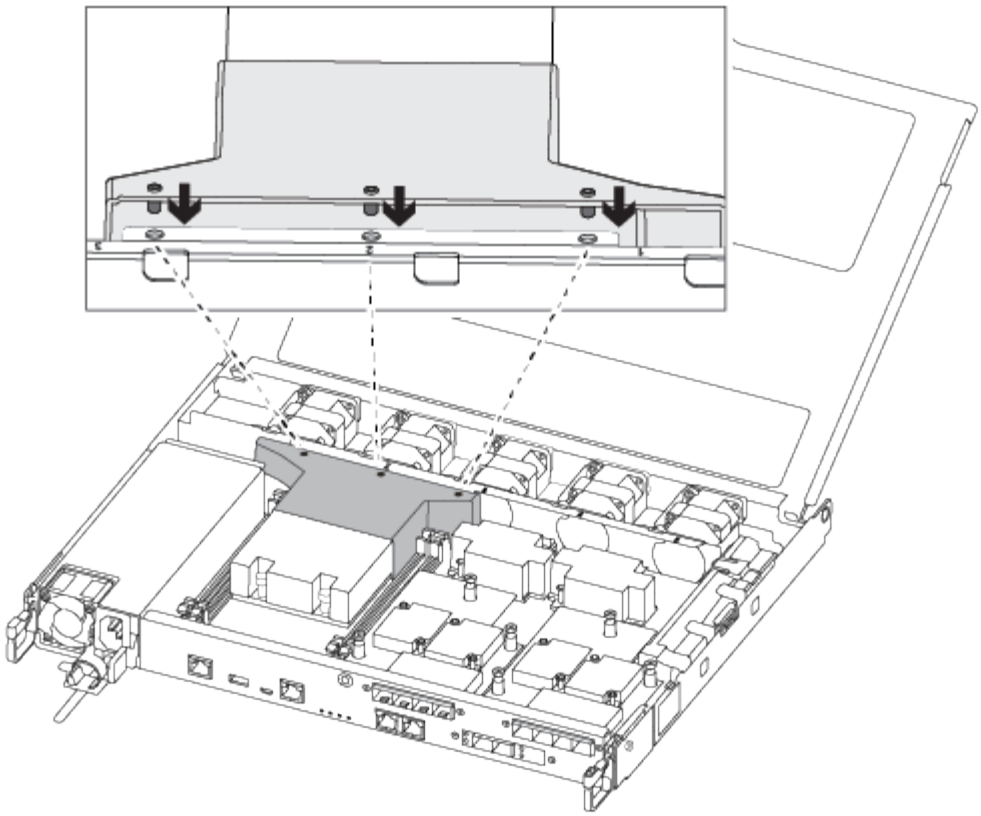

7. Close the controller module cover and tighten the thumbscrew.

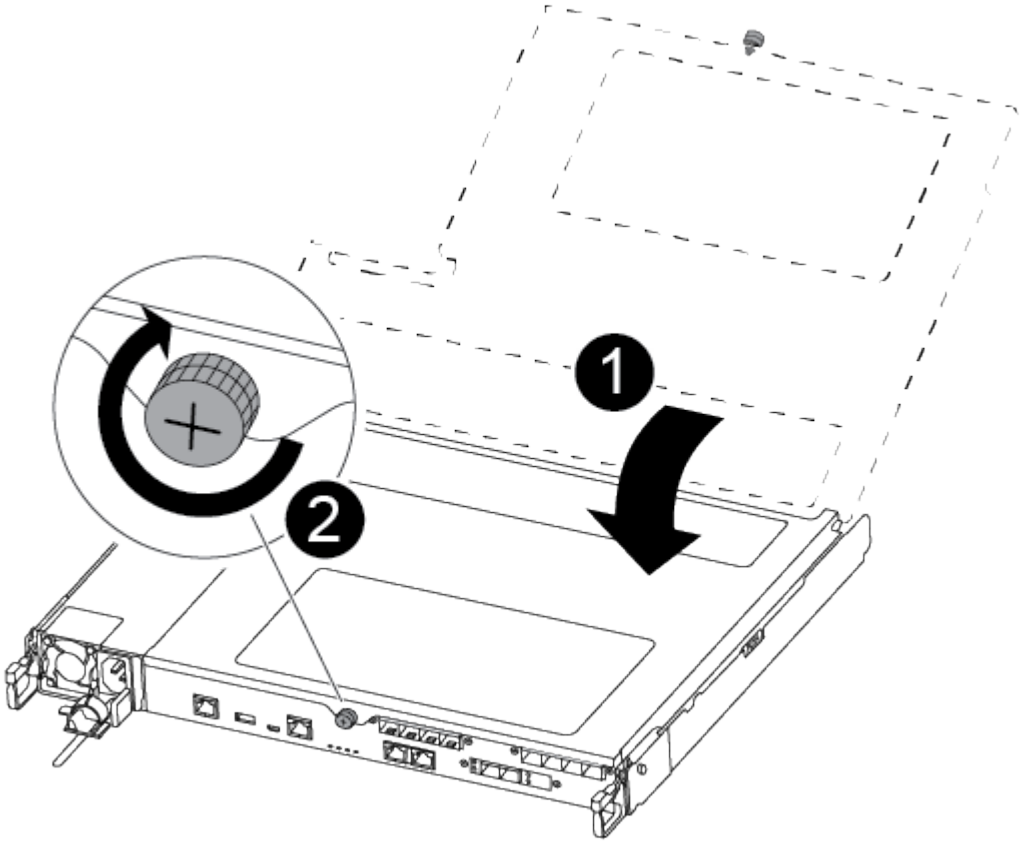

1	Controller module cover
2	Thumbscrew

8. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
9. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

10. Push the controller module all the way into the chassis:
11. Place your index fingers through the finger holes from the inside of the latching mechanism.
12. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
13. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

    The controller module should be fully inserted and flush with the edges of the chassis.

14. Reconnect the controller module I/O cables.
15. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

    The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

16. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

    If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

17. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

    The system begins to boot and stops at the LOADER prompt.

Boot the recovery image - AFF C250

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.

3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.

5. Give back the controller using the `storage failover giveback -fromnode local` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

   NOTE: If the process fails, contact NetApp Support.

Restore encryption - AFF C250

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
**********************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
**********************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF C250

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF C250

To replace the chassis, you must move the bezel, controller modules, and NVMe drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in its place.

About this task

- All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF C250

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

```
Warning: Are you sure you want to halt node <node_name>? {y|n}:
```

10. Wait for each controller to halt and display the LOADER prompt.

Replace hardware - AFF C250

To replace the chassis, you move the power supplies, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis from with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

Use the following video or the tabulated steps to replace the chassis; it assumes the removal and replacement of the bezel:

Animation - Replace the chassis

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

> If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

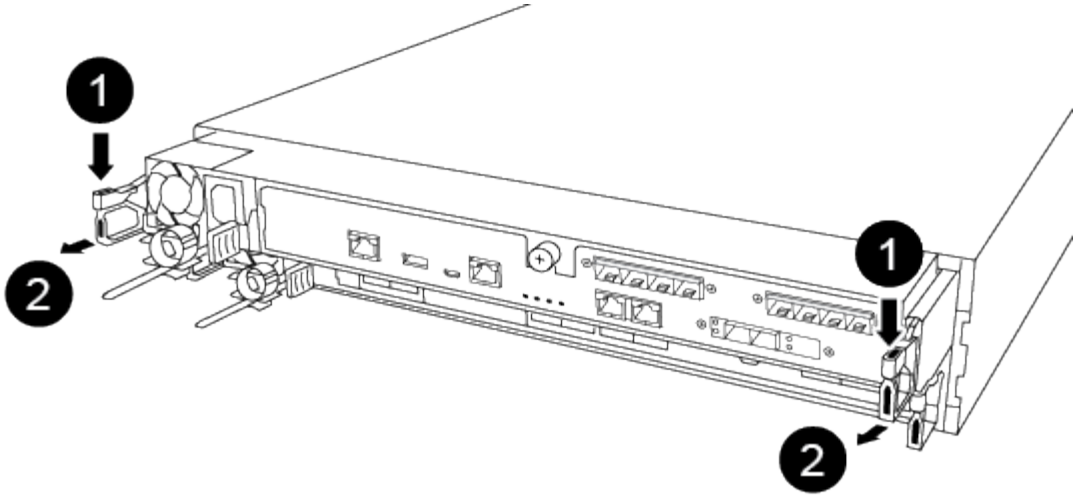

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.
5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up

and against the drive holder.

Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot the system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Insert the controller module into the chassis:

   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

   The controller module should be fully inserted and flush with the edges of the chassis.

   a. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect

the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - AFF C250

You must verify the HA state of the chassis, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis *HA-state*`

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement- AFF C250

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct

recovery procedure to determine whether you should use this procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller module - AFF C250

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - AFF C250

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

Use the following video or the tabulated steps to replace a controller module:

Animation - Replace a controller module

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

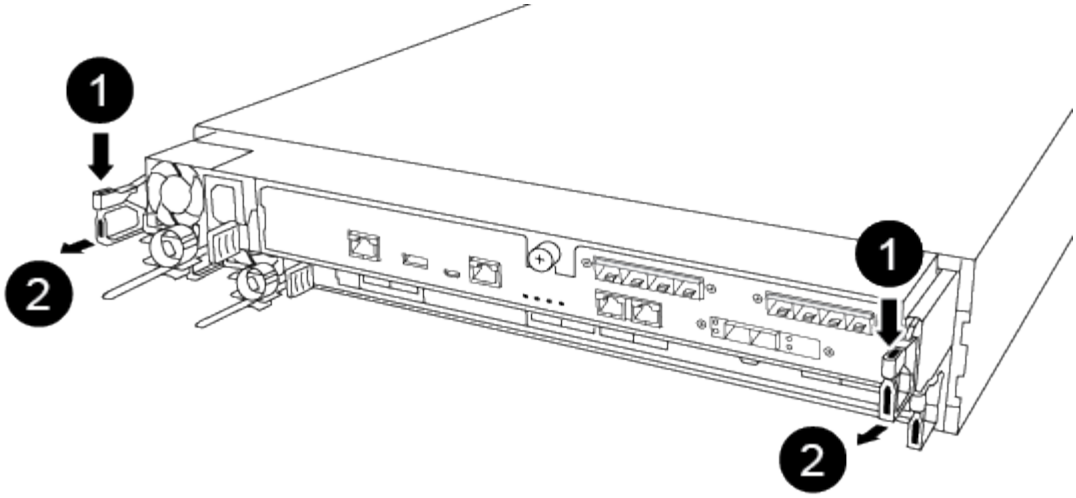

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

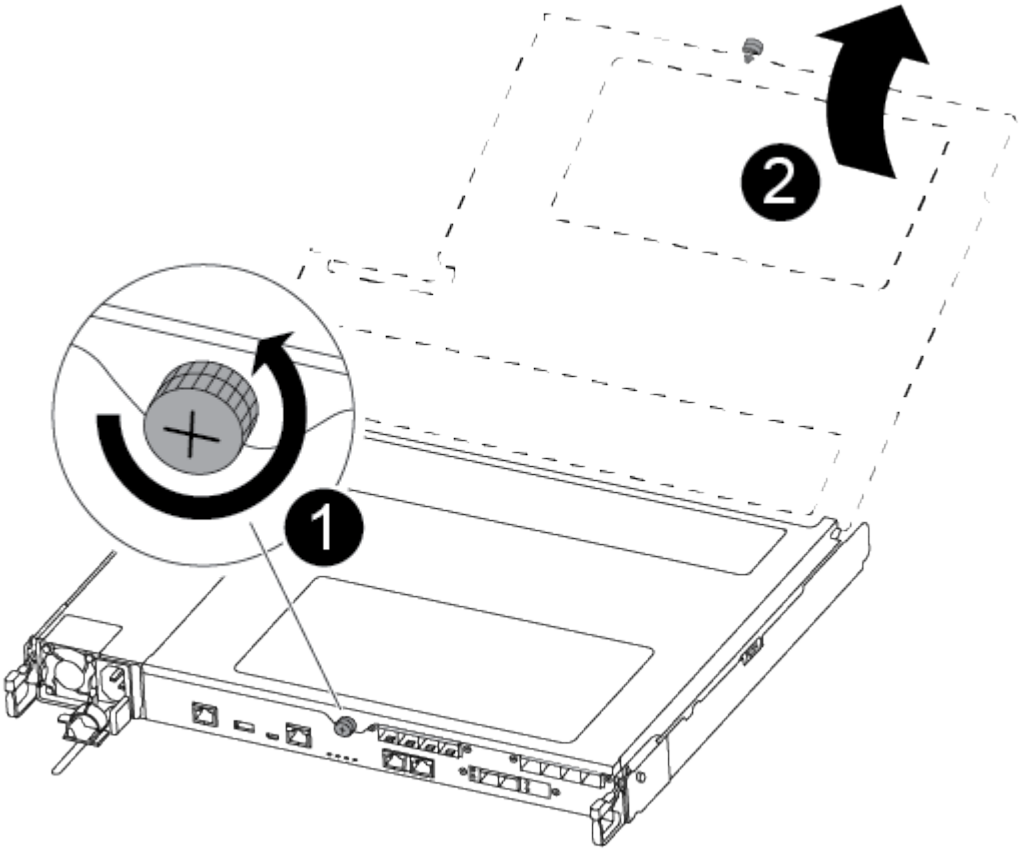

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

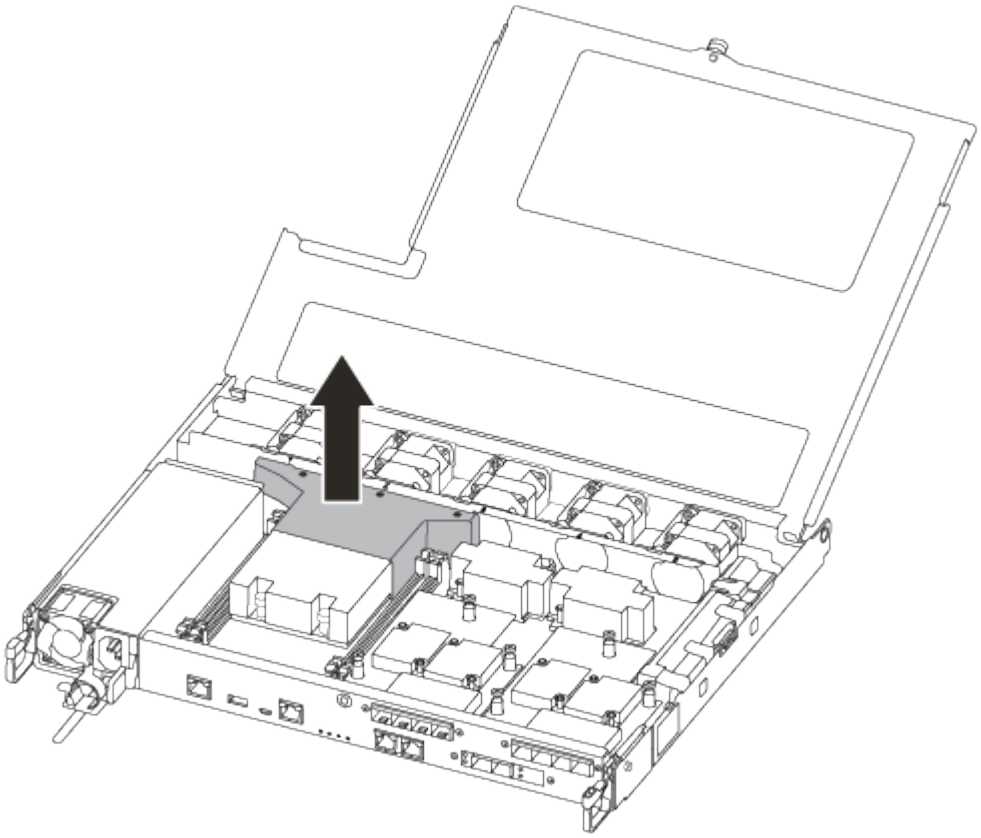

Step 2: Move the power supply

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

1. Disconnect the power supply.
2. Open the power cable retainer, and then unplug the power cable from the power supply.
3. Unplug the power cable from the power source.
4. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

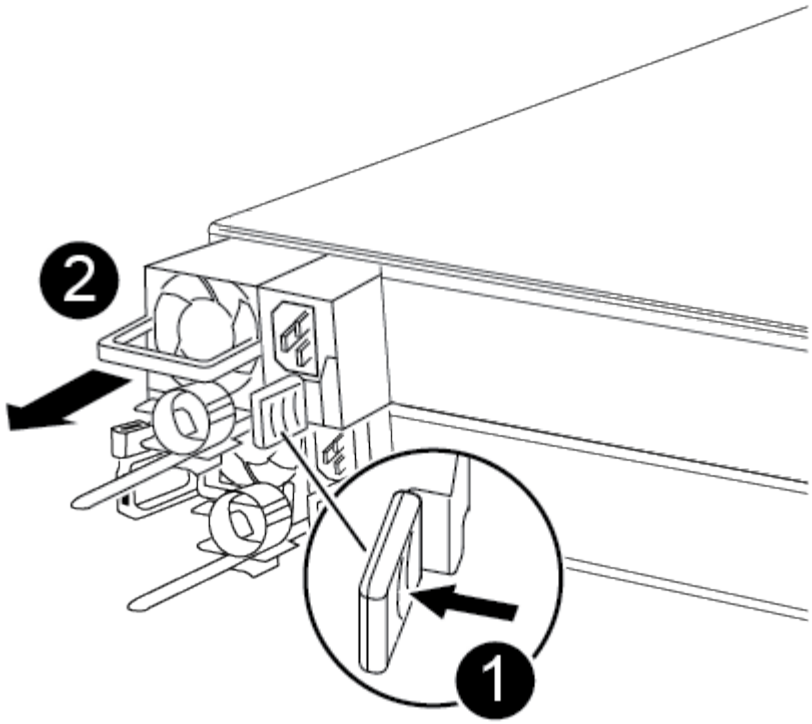

1	Blue power supply locking tab
2	Power supply

5. Move the power supply to the new controller module, and then install it.
6. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

You must move the fans from the impaired controller module to the replacement module when replacing a failed controller module.

1. Remove the fan module by pinching the side of the fan module, and then lifting the fan module straight out of the controller module.

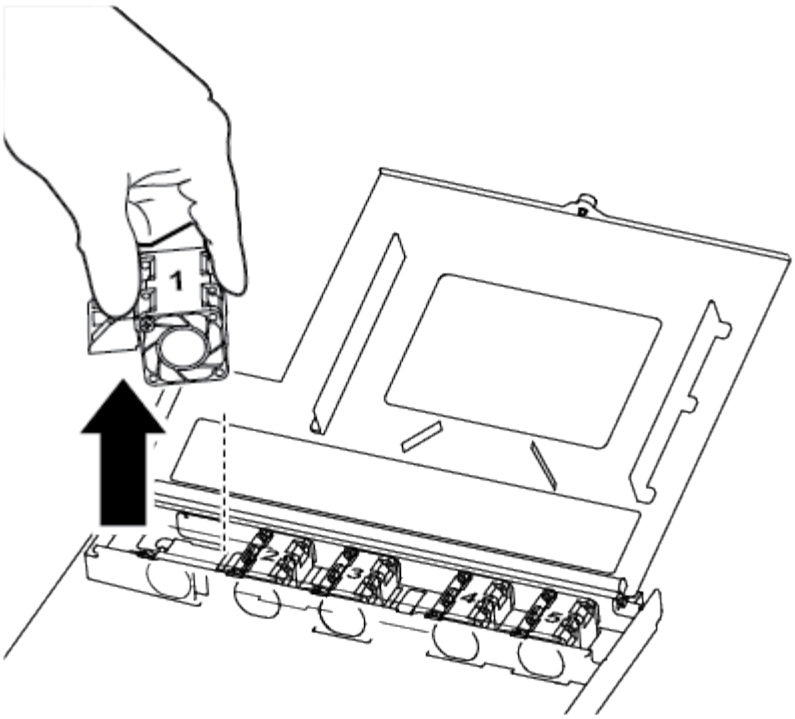

1	Fan module

2. Move the fan module to the replacement controller module, and align the edges of the fan module with the opening in the controller module, and then slide the fan module in.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the boot media

You must move the boot media device from the impaired controller module to the replacement controller module.

You need a #1 magnetic Phillips head screwdriver to remove the screw that holds the boot media in place. Due to the space constraints within the controller module, you should also have a magnet to transfer the screw on to so that you do not lose it.

1. Locate and move the boot media from the impaired controller module to the replacement controller module.

   The boot media is located under the air duct cover you removed earlier in this procedure.

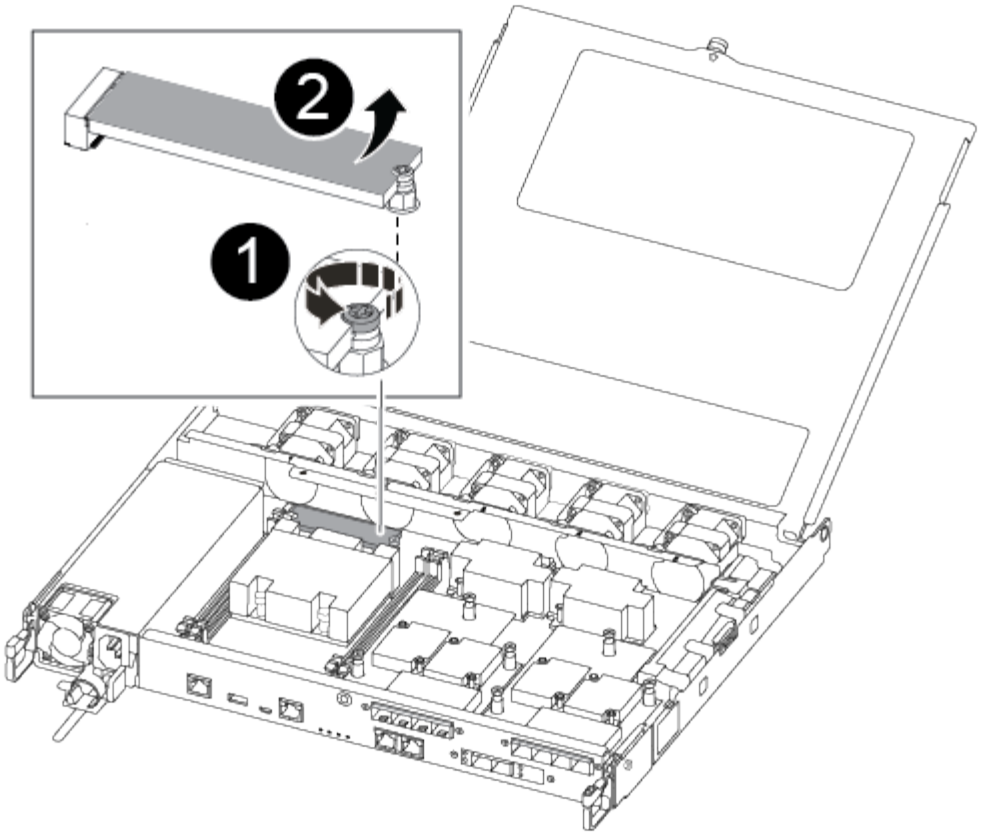

1	Remove the screw securing the boot media to the motherboard in the impaired controller module.
2	Lift the boot media out of the impaired controller module.

2. Using the #1 magnetic screwdriver, remove the screw from the boot media, and set it aside safely on the magnet.
3. Gently lift the boot media directly out of the socket and align it into place in the replacement controller module.
4. Using the #1 magnetic screwdriver, insert and tighten the screw on the boot media.

> Do not apply force when tightening the screw on the boot media; you might crack it.

Step 5: Move the DIMMs

To move the DIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

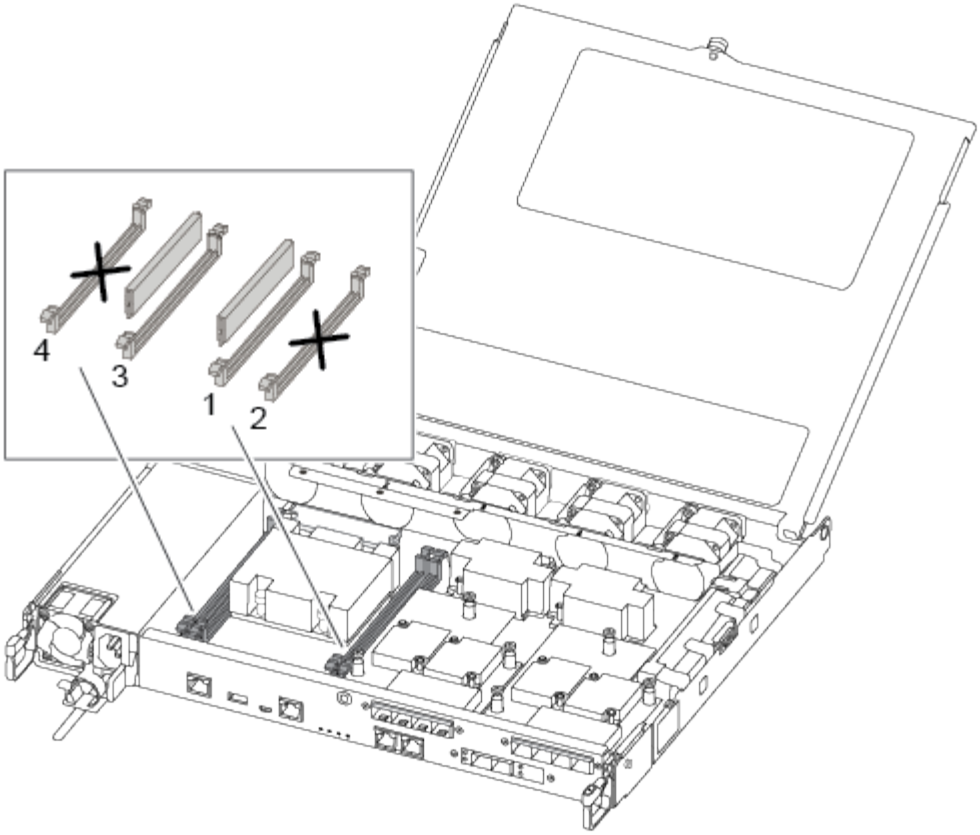

Install each DIMM into the same slot it occupied in the impaired controller module.

1. Slowly push apart the DIMM ejector tabs on either side of the DIMM, and slide the DIMM out of the slot.

   Hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

2. Locate the corresponding DIMM slot on the replacement controller module.
3. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the DIMM squarely into the socket.

   The DIMMs fit tightly in the socket. If not, reinsert the DIMM to realign it with the socket.

4. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.
5. Repeat these steps for the remaining DIMM.

Step 6: Move a mezzanine card

To move a mezzanine card, you must remove the cabling and any QSFPs and SFPs from the ports, move the mezzanine card to the replacement controller, reinstall any QSFPs and SFPs onto the ports, and cable the ports.

1. Locate and move the mezzanine cards from your impaired controller module.

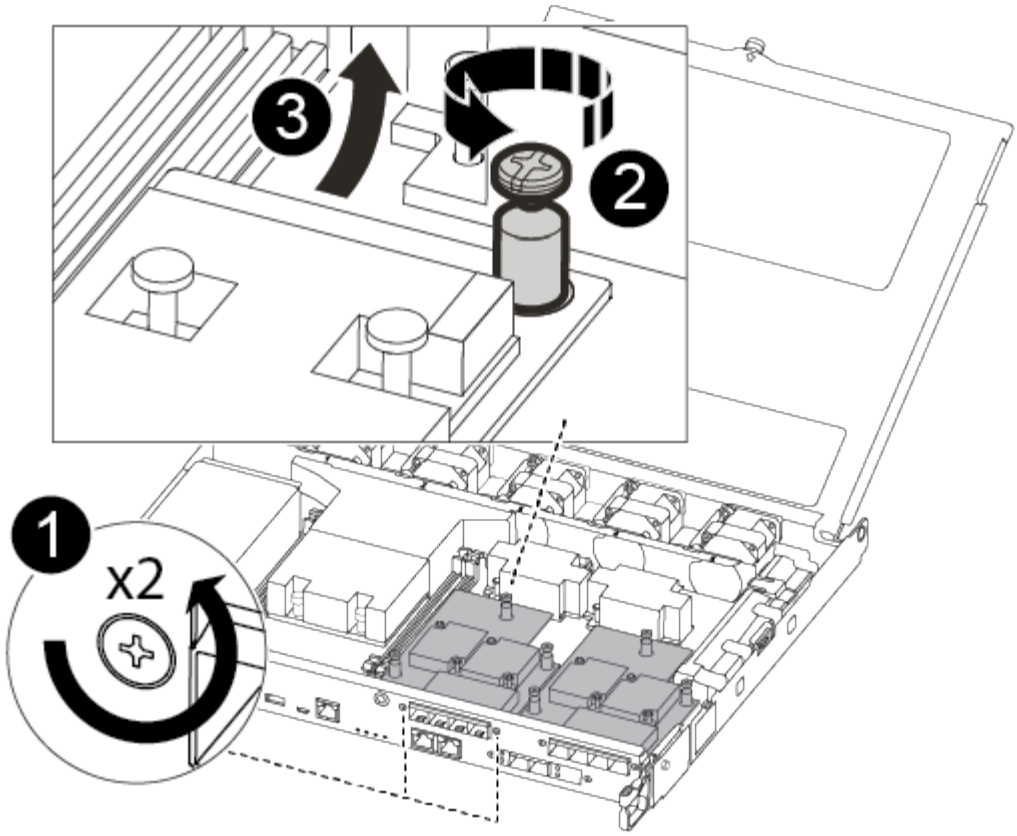

1	Remove screws on the face of the controller module.
2	Loosen the screw in the controller module.
3	Move the mezzanine card.

2. Unplug any cabling associated with the mezzanine card.

   Make sure that you label the cables so that you know where they came from.

   a. Remove any SFP or QSFP modules that might be in the mezzanine card and set it aside.
   b. Using the #1 magnetic screwdriver, remove the screws from the face of the impaired controller module and from the mezzanine card, and set them aside safely on the magnet.
   c. Gently lift the mezzanine card out of the socket and move it to the same position in the replacement controller.
   d. Gently align the mezzanine card into place in the replacement controller.
   e. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the replacement controller module and on the mezzanine card.

   ⓘ Do not apply force when tightening the screw on the mezzanine card; you might crack it.

3. Repeat these steps if there is another mezzanine card in the impaired controller module.

4. Insert the SFP or QSFP modules that were removed onto the mezzanine card.

Step 7: Move the NV battery

When replacing the controller module, you must move the NV battery from the impaired controller module to the replacement controller module.

1. Locate and move the NVMEM battery from your impaired controller module to the replacement controller module.

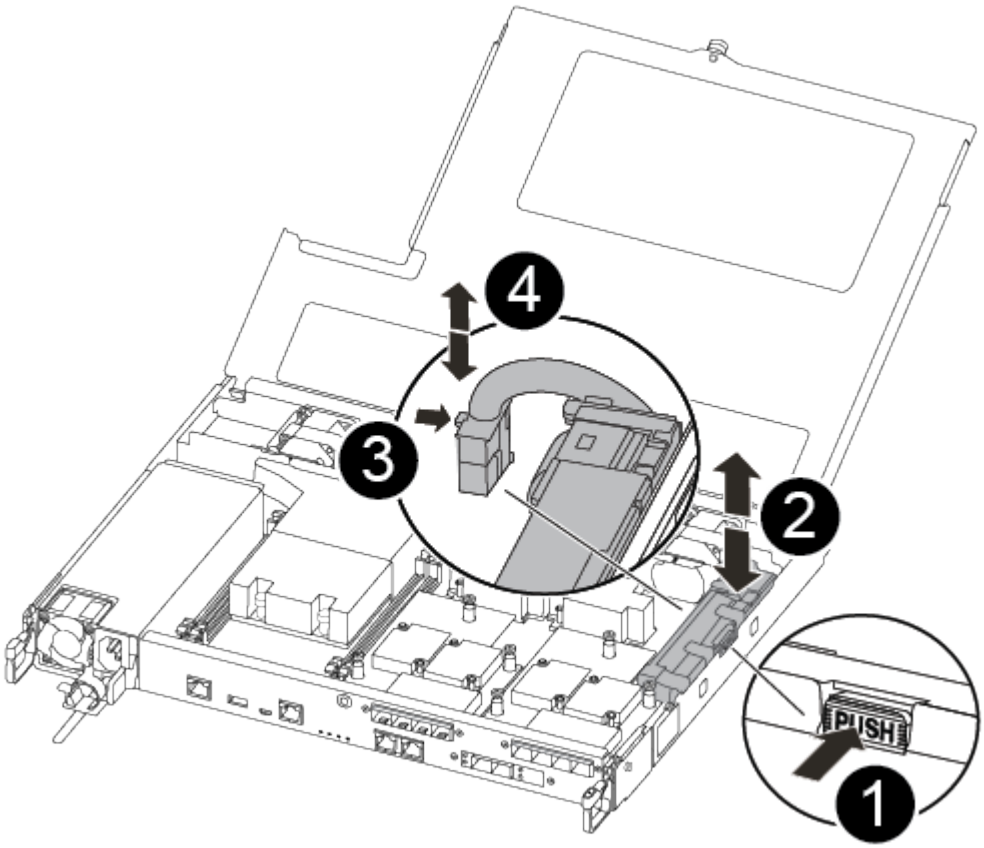

1	Squeeze the clip on the face of the battery plug.
2	Unplug the battery cable from the socket.
3	Grasp the battery and press the blue locking tab marked PUSH.
4	Lift the battery out of the holder and controller module.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.

4. Locate the corresponding NV battery holder on the replacement controller module and align the NV battery to the battery holder.
5. Insert the NV battery plug into the socket.
6. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
7. Press firmly down on the battery pack to make sure that it is locked into place.

Step 8: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

You can use the following illustrations or the written steps to install the replacement controller module in the chassis.

1. If you have not already done so, install the air duct.

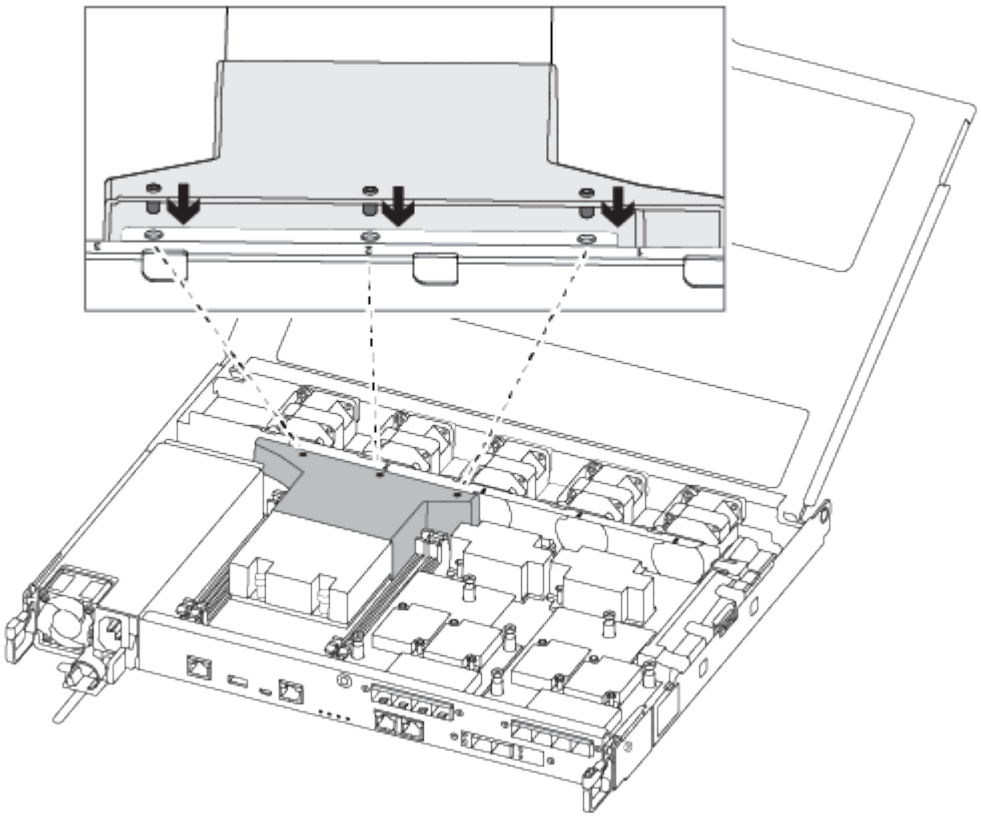

2. Close the controller module cover and tighten the thumbscrew.

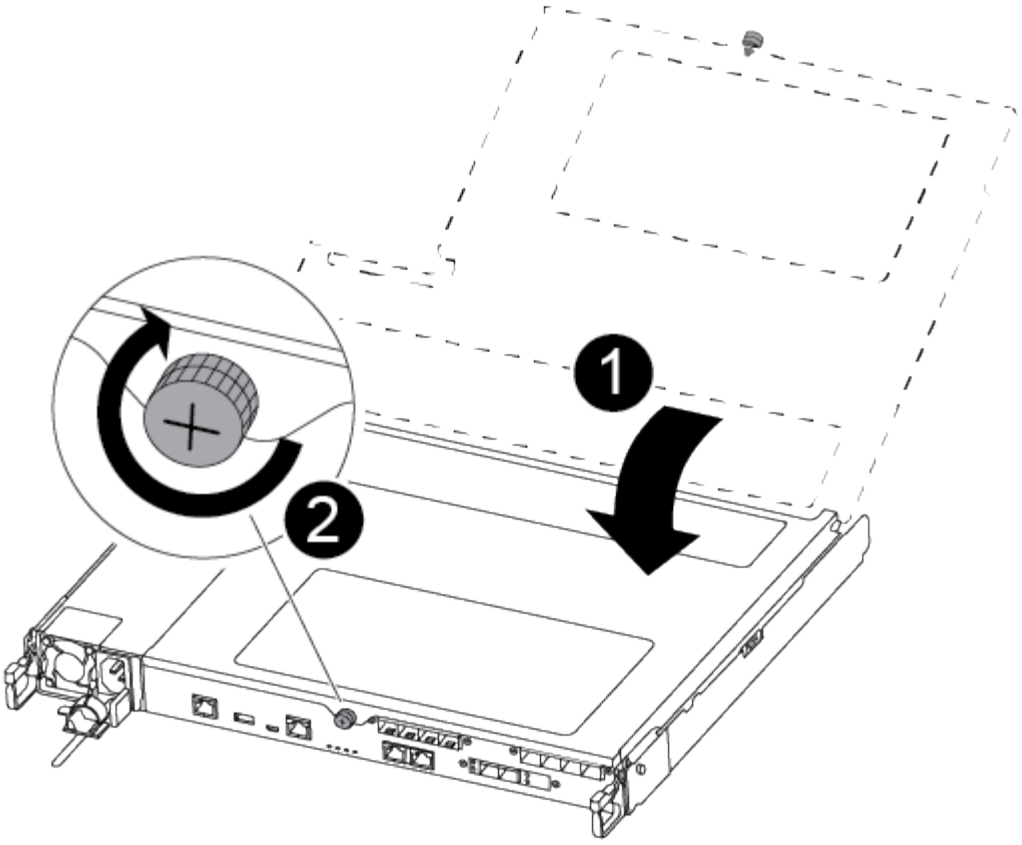

1	Controller module cover
2	Thumbscrew

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching

mechanisms snap into place.

The controller module should be fully inserted and flush with the edges of the chassis.

f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

Restore and verify the system configuration - AFF C250

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - ha
   - mcc
   - mccip
   - non-ha

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF C250

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and

then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF C250

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF C250

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

About this task

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

> If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

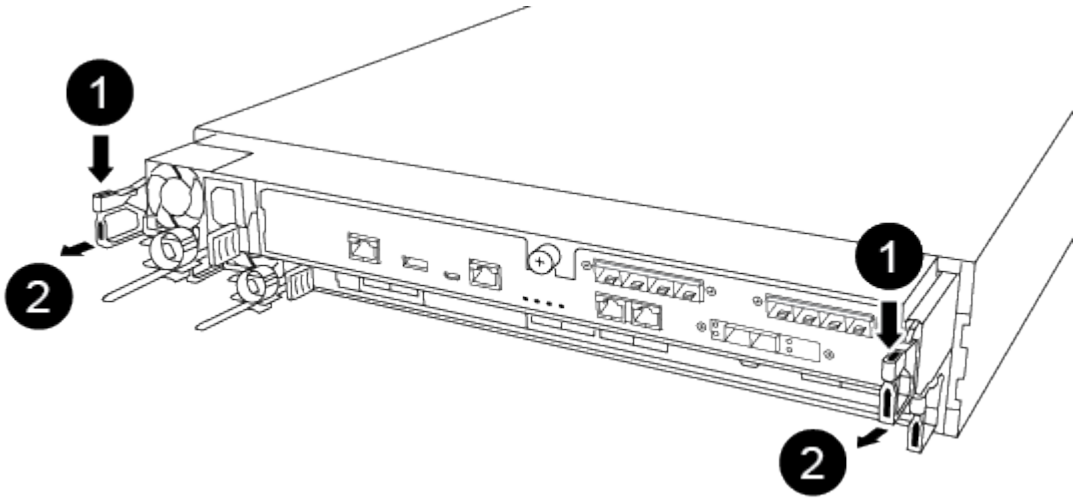

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

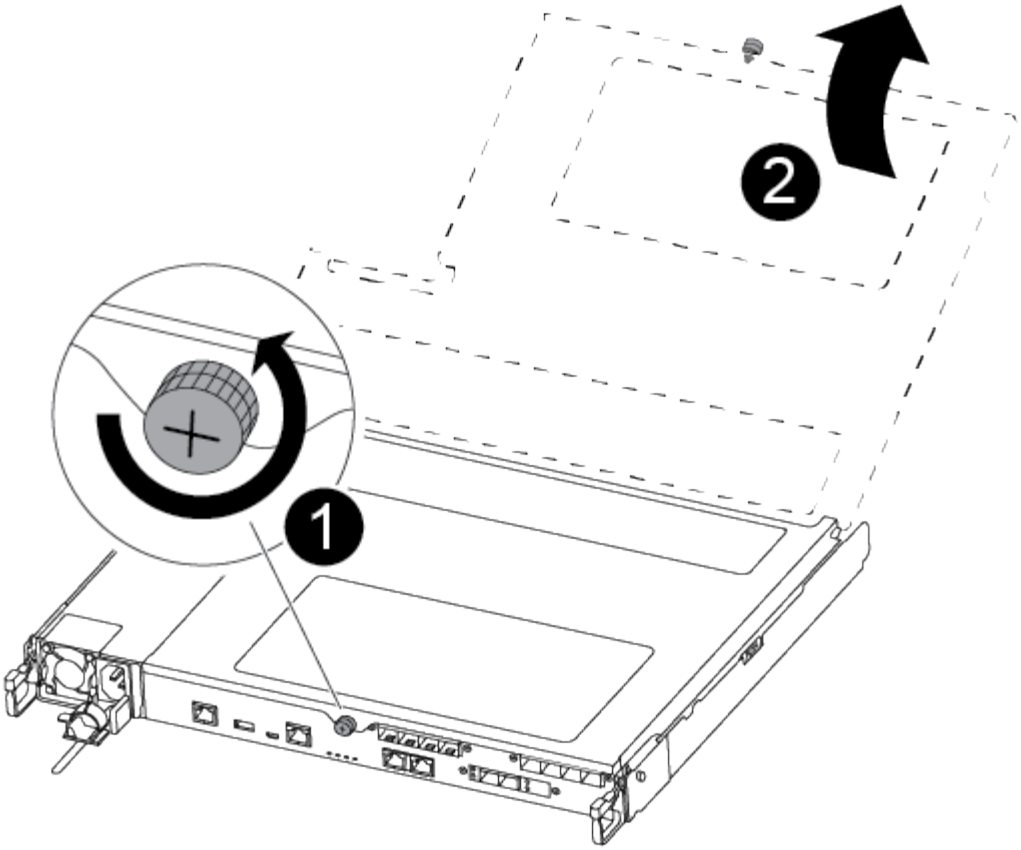

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

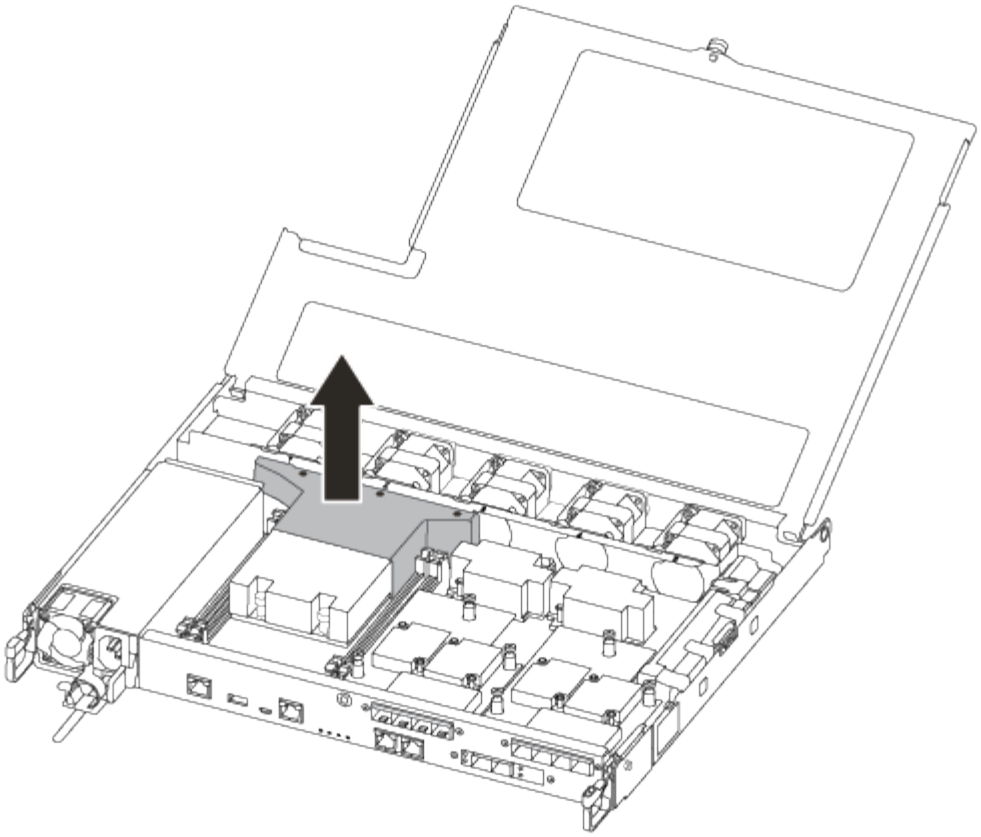

Step 3: Replace a DIMM

To replace a DIMM, you must locate it in the controller module using the DIMM map label on top of the air duct and then replace it following the specific sequence of steps.

Use the following video or the tabulated steps to replace a DIMM:

Animation - Replace a DIMM

1. Replace the impaired DIMM on your controller module.

   The DIMMs are in slot 3 or 1 on the motherboard. Slot 2 and 4 are left empty. Do not attempt to install DIMMs into these slots.

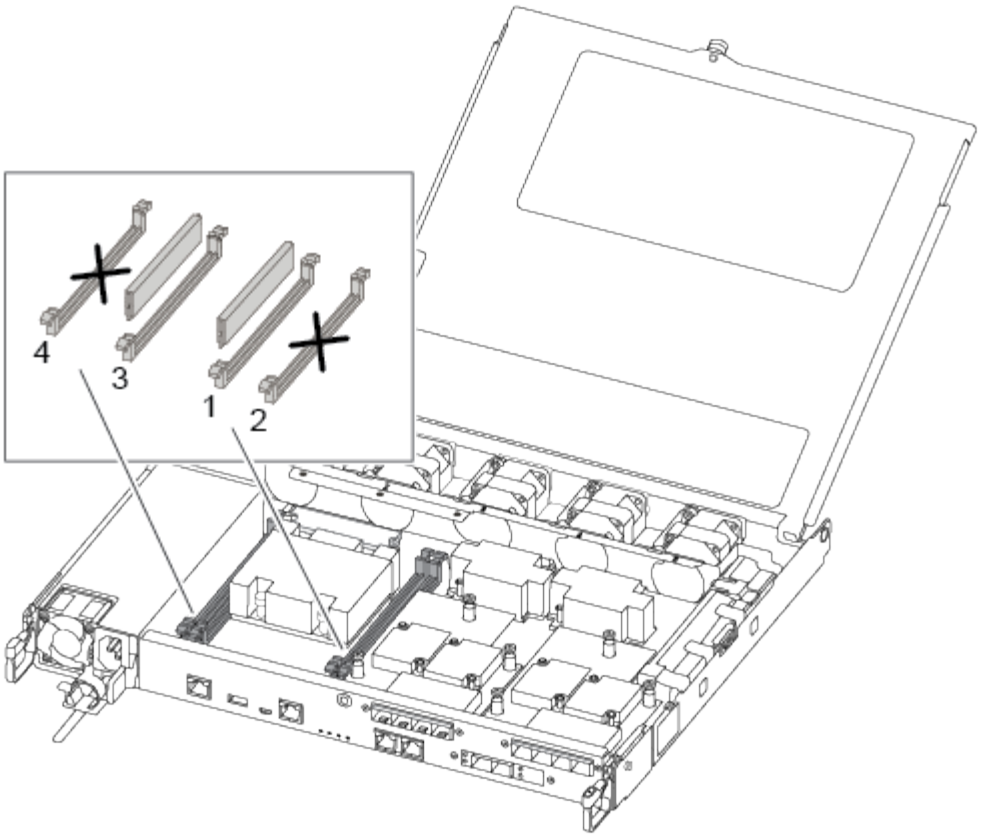

2. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
3. Slowly push apart the DIMM ejector tabs on either side of the DIMM, and slide the DIMM out of the slot.
4. Leave DIMM ejector tabs on the connector in the open position.
5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   Hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

6. Insert the replacement DIMM squarely into the slot.

   The DIMMs fit tightly in the socket. If not, reinsert the DIMM to realign it with the socket.

7. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following illustrations or the written steps to install the replacement controller module in the chassis.

1. If you have not already done so, install the air duct.

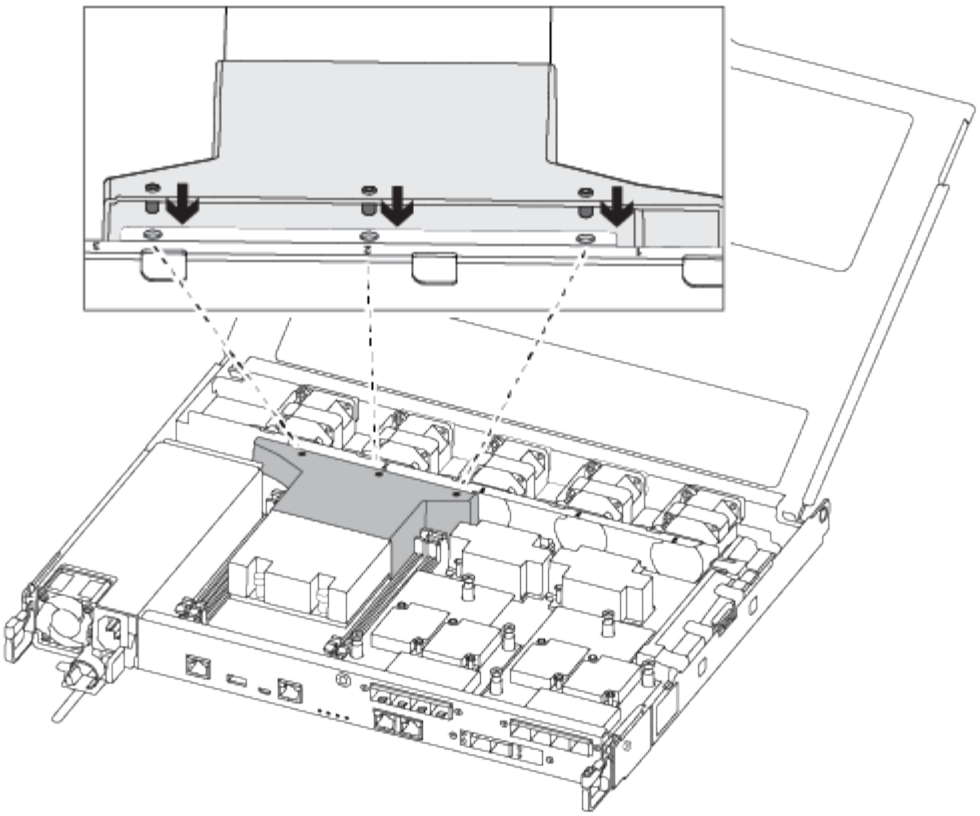

2. Close the controller module cover and tighten the thumbscrew.

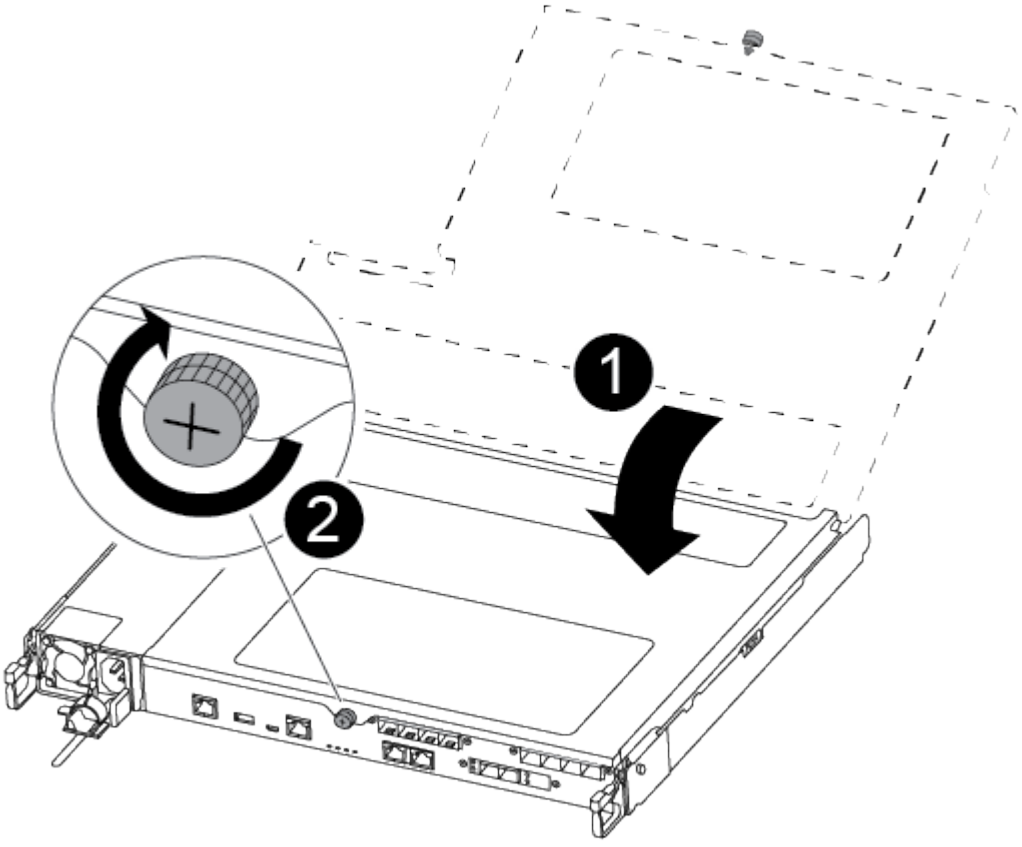

1	Controller module cover
2	Thumbscrew

3. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.
4. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.
5. Recable the system, as needed.
6. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
7. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD drive - AFF C250

You can replace a failed SSD drive nondisruptively while I/O is in progress.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed disk drive by running the `storage disk show -broken` command from the system console.

The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

It can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the disk depends on how the disk drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

When replacing several disk drives, you must wait 70 seconds between the removal of each failed disk drive and the insertion of the replacement disk drive to allow the storage system to recognize the existence of each new disk.

You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   > The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

4. Remove the failed drive:

a. Press the release button on the drive face to open the cam handle.
b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.

5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

7. Verify that the drive's activity (green) LED is illuminated.

   When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      ```
      storage disk show -container-type unassigned
      ```

      You can enter the command on either controller module.

   b. Assign each drive:

      ```
      storage disk assign -disk disk_name -owner node_name
      ```

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      ```
      storage disk option modify -node node_name -autoassign on
      ```

      You must reenable automatic drive assignment on both controller modules.

10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Replace a fan - AFF C250

To replace a fan, remove the failed fan module and replace it with a new fan module.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

> If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

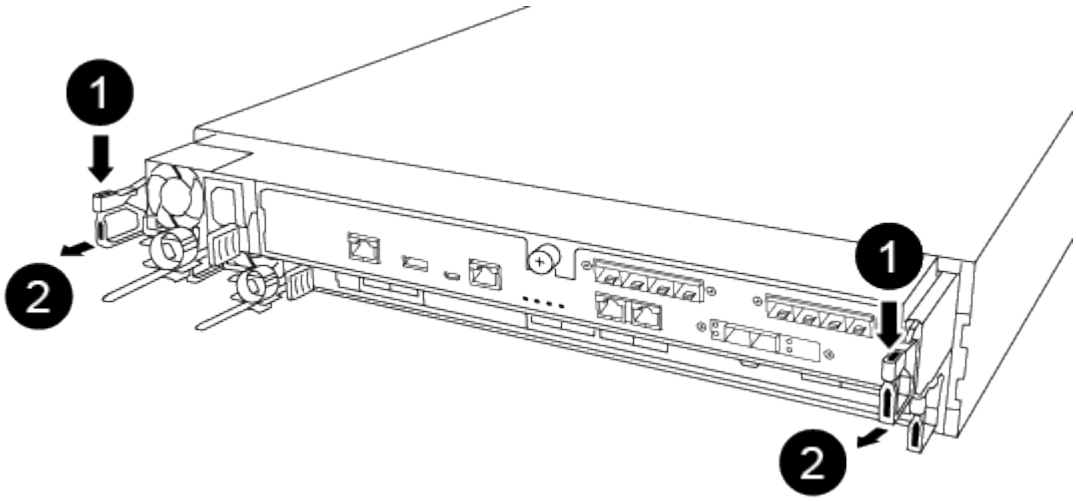

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

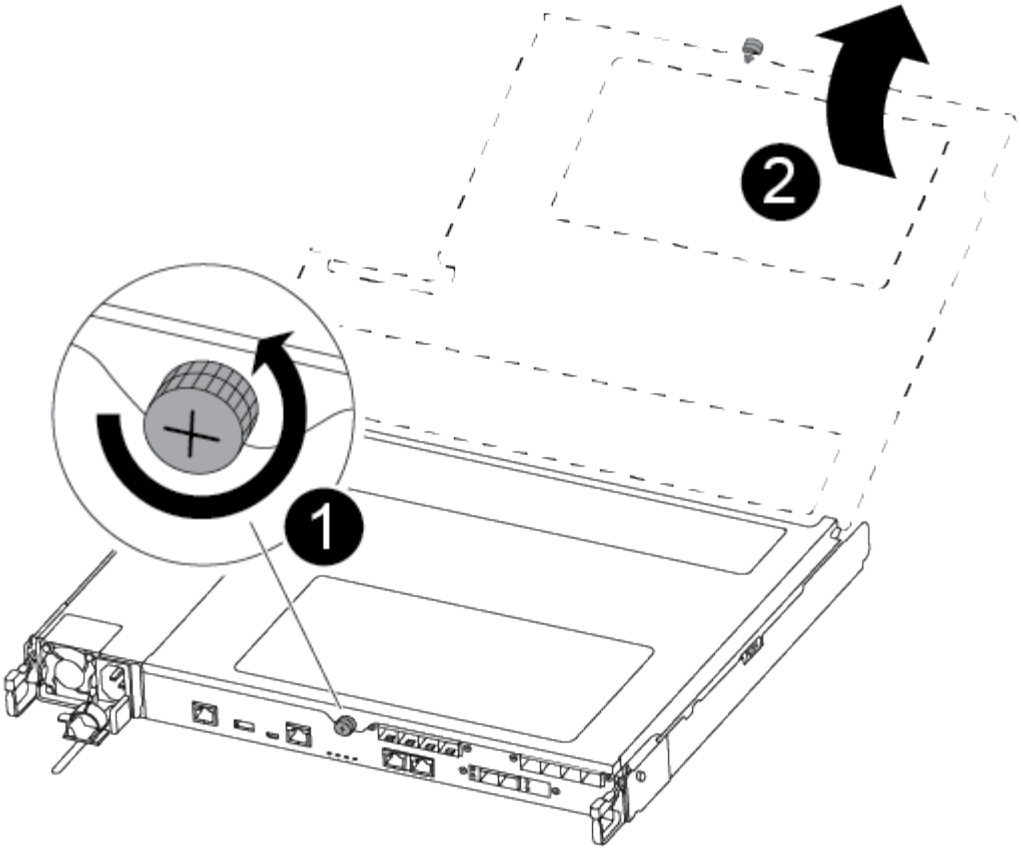

1	Thumbscrew
2	Controller module cover

Step 3: Replace a fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

Use the following video or the tabulated steps to replace a fan:

Animation - Replace a fan

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
2. Remove the fan module by pinching the side of the fan module, and then lifting the fan module straight out of the controller module.

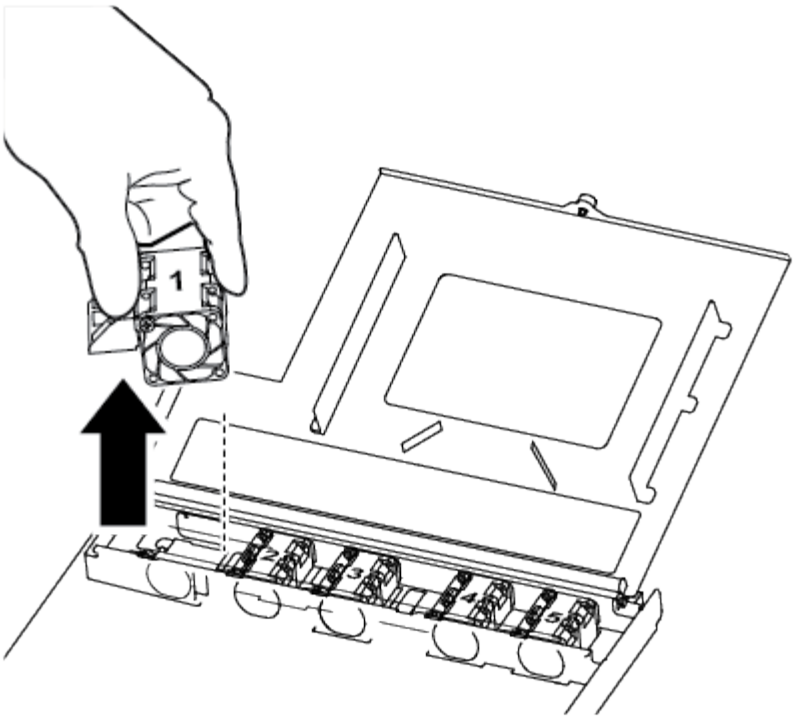

1	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Close the controller module cover and tighten the thumbscrew.

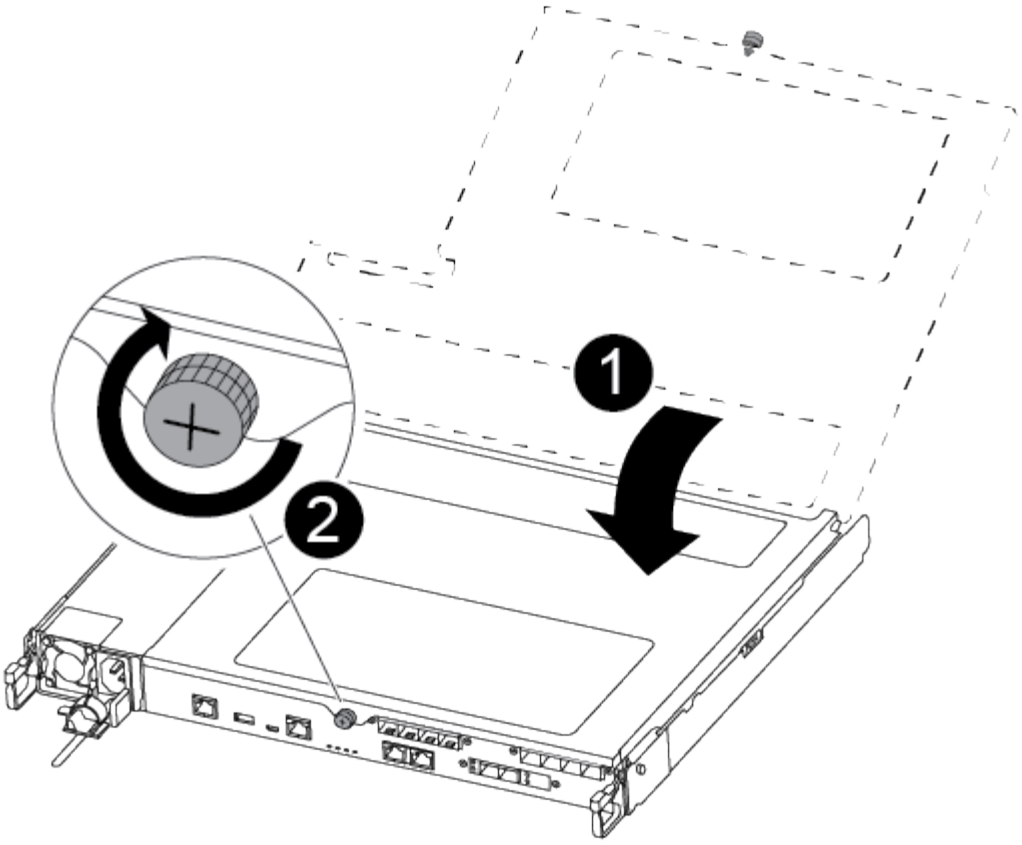

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.

   f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace or install a mezzanine card - AFF C250

To replace a failed mezzanine card, you must remove the cables and any SFP or QSFP modules, replace the card, reinstall the SFP or QSFP modules and recable the cards. To install a new mezzanine card, you must have the appropriate cables and SFP or QSFP modules.

About this task

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

Remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

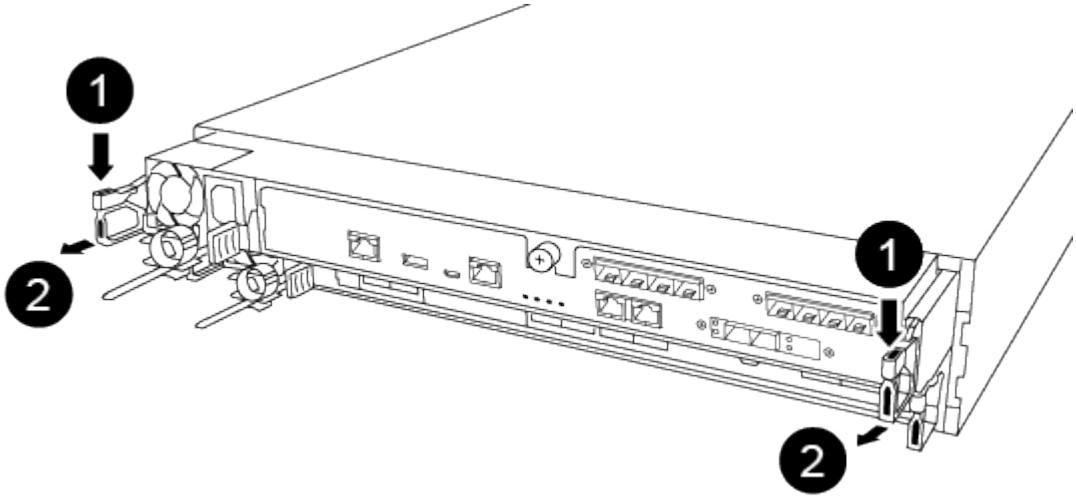

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

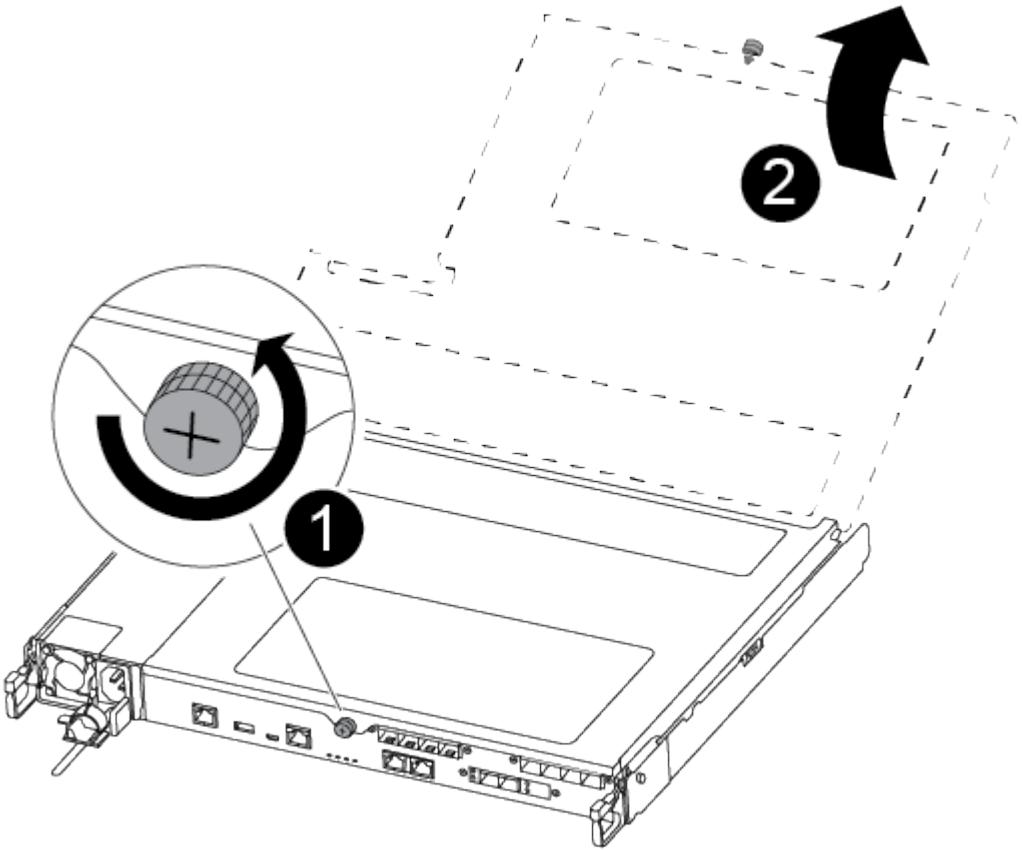

1	Thumbscrew
2	Controller module cover.

Step 3: Replace or install a mezzanine card

To replace a mezzanine card, you must remove the impaired card and install the replacement card; to install a mezzanine card, you must remove the faceplate and install the new card.

Use the following video or the tabulated steps to replace a mezzanine card:

Animation - Replace a mezzanine card

1. To replace a mezzanine card:
2. Locate and replace the impaired mezzanine card on your controller module.

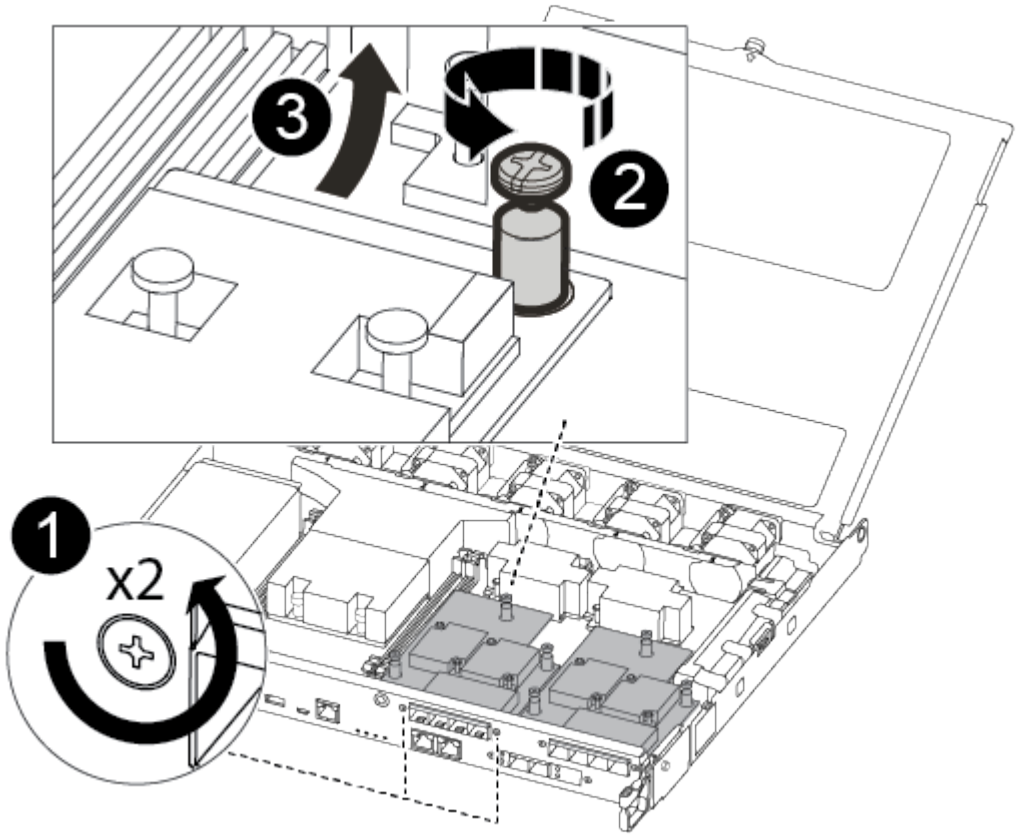

1	Remove screws on the face of the controller module.
2	Loosen the screw in the controller module.
3	Remove the mezzanine card.

a. Unplug any cabling associated with the impaired mezzanine card.

   Make sure that you label the cables so that you know where they came from.

b. Remove any SFP or QSFP modules that might be in the impaired mezzanine card and set it aside.
c. Using the #1 magnetic screwdriver, remove the screws from the face of the controller module and set them aside safely on the magnet.
d. Using the #1 magnetic screwdriver, loosen the screw on the impaired mezzanine card.
e. Using the #1 magnetic screwdriver, gently lift the impaired mezzanine card directly out of the socket and set it aside.
f. Remove the replacement mezzanine card from the antistatic shipping bag and align it to the inside face of the controller module.
g. Gently align the replacement mezzanine card into place.

h. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the controller module and on the mezzanine card.

> Do not apply force when tightening the screw on the mezzanine card; you might crack it.

i. Insert any SFP or QSFP modules that were removed from the impaired mezzanine card to the replacement mezzanine card.

3. To install a mezzanine card:
4. You install a new mezzanine card if your system does not have one.
   a. Using the #1 magnetic screwdriver, remove the screws from the face of the controller module and the faceplate covering the mezzanine card slot, and set them aside safely on the magnet.
   b. Remove the mezzanine card from the antistatic shipping bag and align it to the inside face of the controller module.
   c. Gently align the mezzanine card into place.
   d. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the controller module and on the mezzanine card.

> Do not apply force when tightening the screw on the mezzanine card; you might crack it.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Close the controller module cover and tighten the thumbscrew.

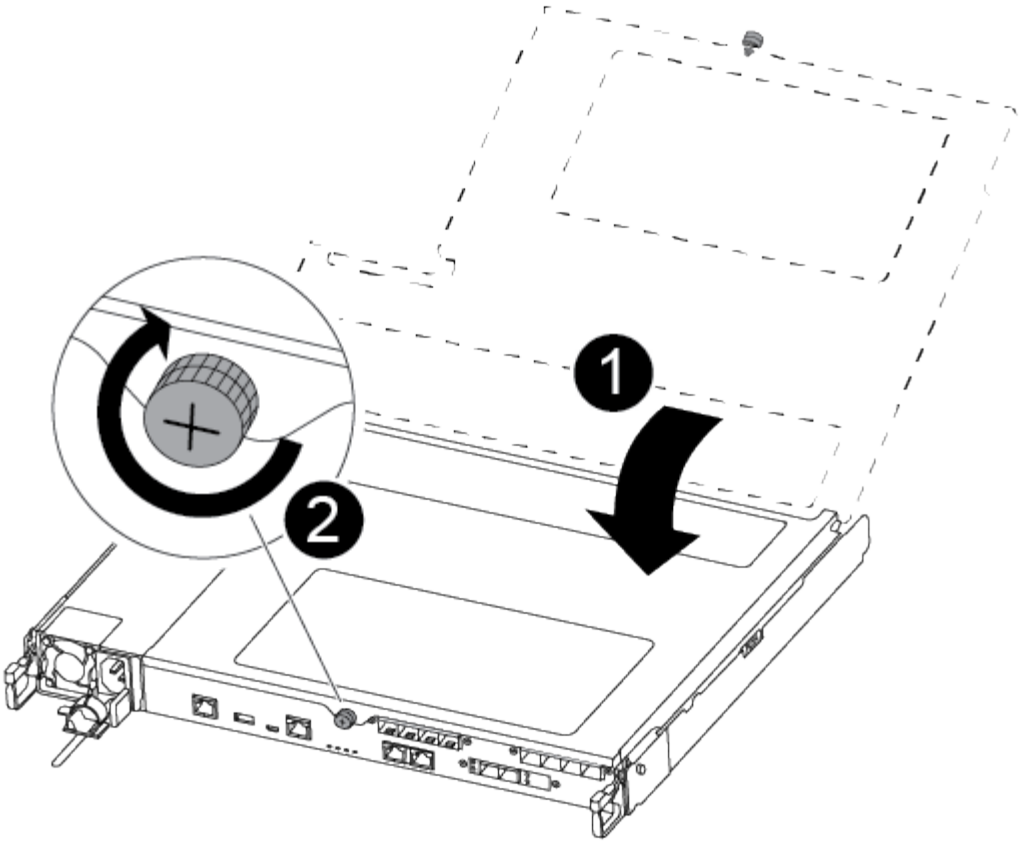

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.

   f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVMEM battery - AFF C250

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

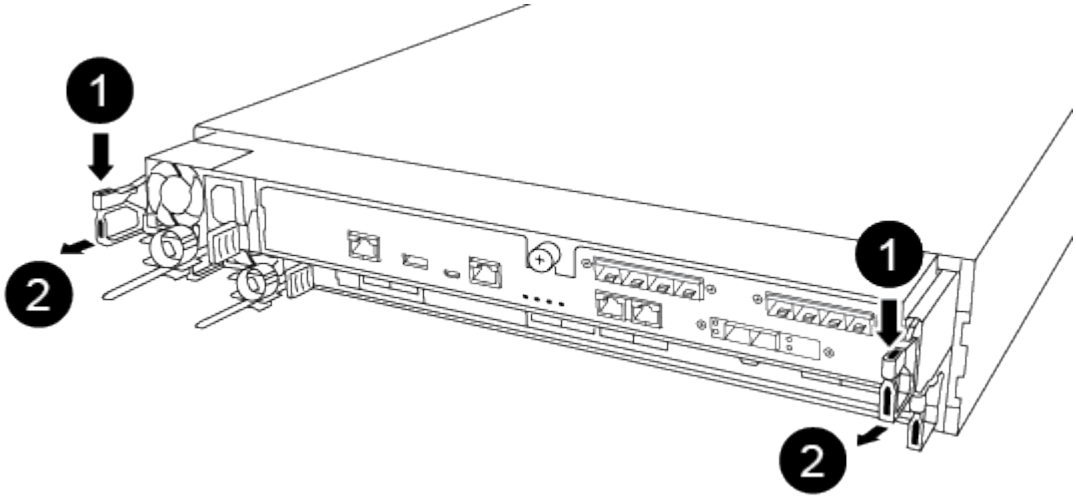

1	Lever

2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

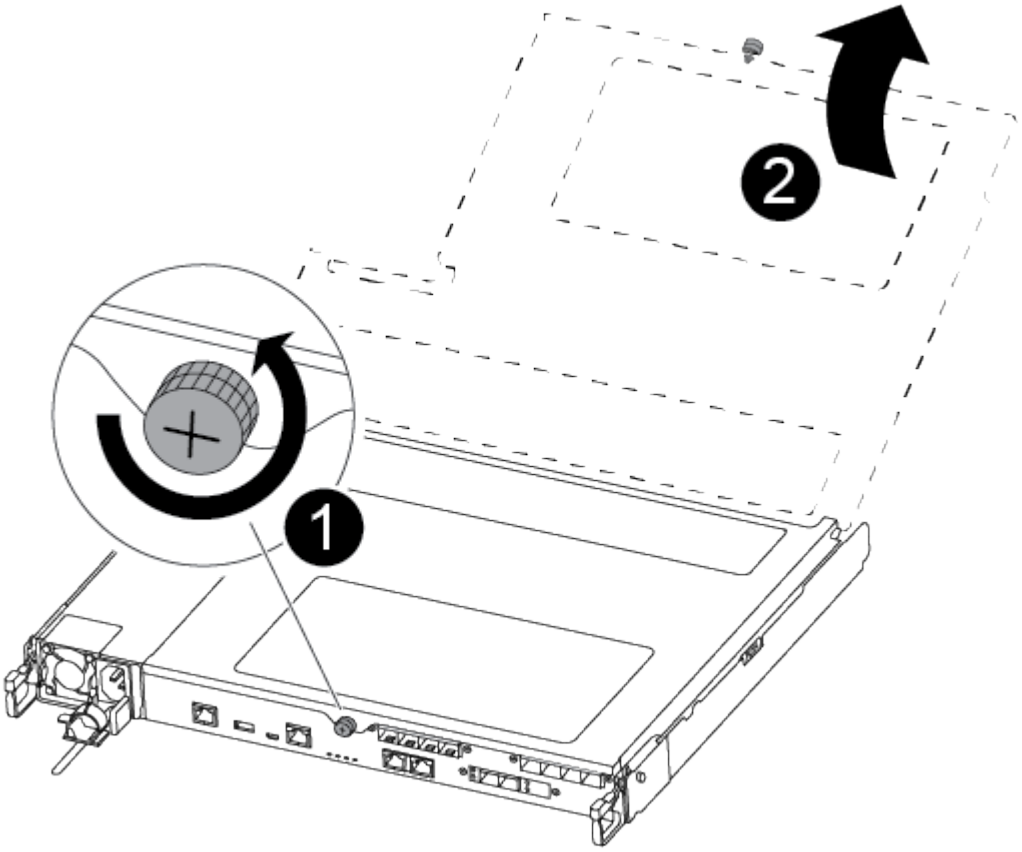

1	Thumbscrew
2	Controller module cover.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module.

Use the following video or the tabulated steps to replace the NVMEM battery:

Animation - Replace the NVMEM battery

1. Locate and replace the impaired NVMEM battery on your controller module.

It is recommended that you follow the illustrated instructions in the order listed.

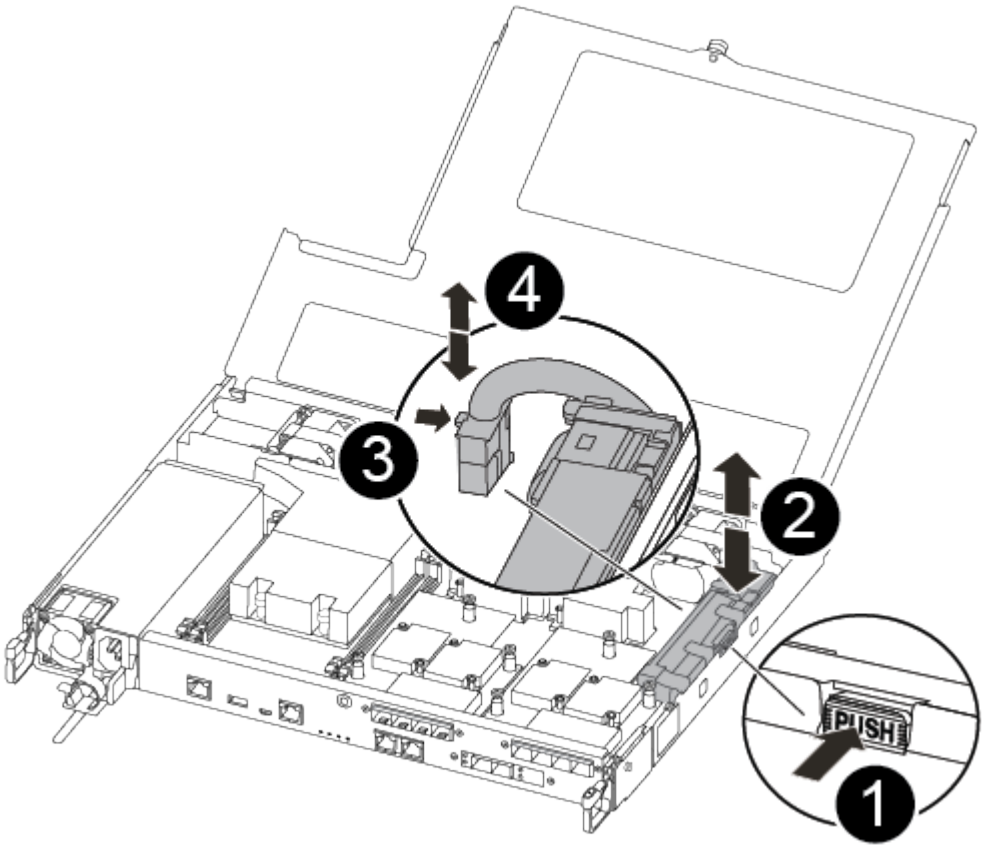

1	Squeeze the clip on the face of the battery plug.
2	Unplug the battery cable from the socket.
3	Grasp the battery and press the blue locking tab marked PUSH.
4	Lift the battery out of the holder and controller module.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module and set it aside.
4. Remove the replacement NV battery from the antistatic shipping bag and align it to the battery holder.
5. Insert the replacement NV battery plug into the socket.
6. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
7. Press firmly down on the battery pack to make sure that it is locked into place.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following illustration or the written steps to install the replacement controller module in the chassis.

1. Close the controller module cover and tighten the thumbscrew.

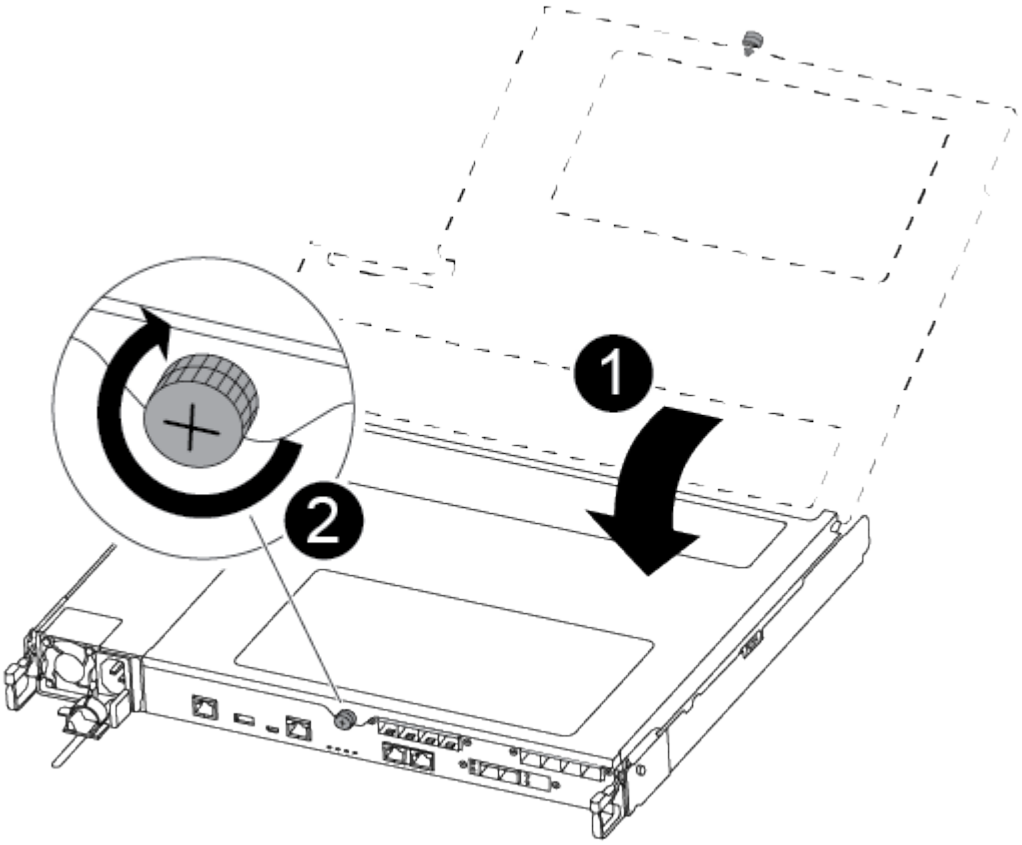

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

The controller module should be fully inserted and flush with the edges of the chassis.

f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.
4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - AFF C250

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one PSU at a time.

> It is a best practice to replace the PSU within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded PSU until the PSU is replaced.

- Power supplies are auto-ranging.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

Use the appropriate procedure for your type of PSU; AC or DC.

Option 1: Replace an AC PSU

Use the following video or the tabulated steps to replace the PSU:

Animation - Replace the AC PSU

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

> The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

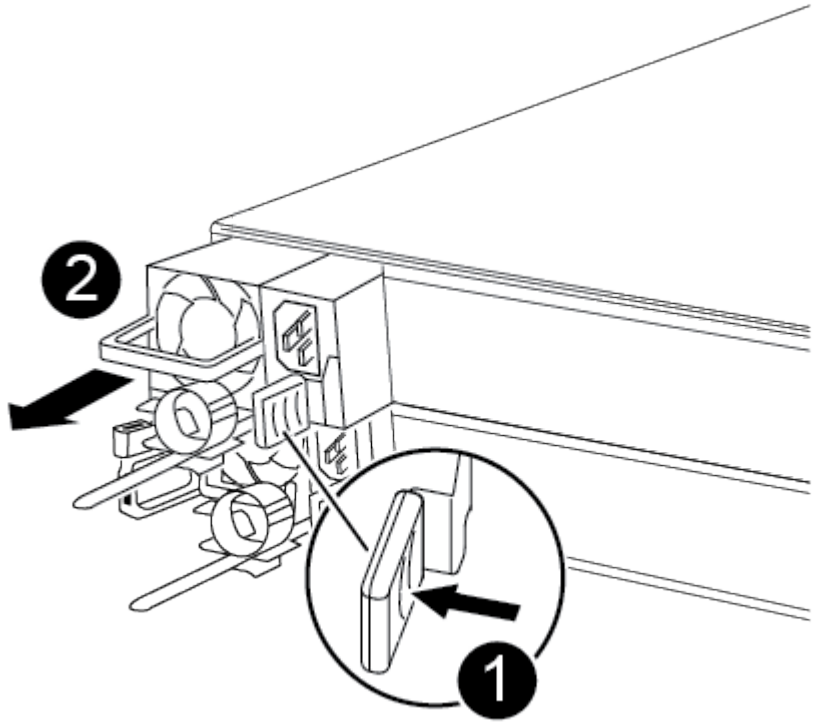

1	Blue PSU locking tab
2	Power supply

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller module.

b. Gently push the PSU into the controller module until the locking tab clicks into place.

The power supplies will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:

a. Reconnect the power cable to the PSU.

b. Secure the power cable to the PSU using the power cable retainer.

Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC power cable connector using the thumb screws on the plug.
   b. Unplug the power cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

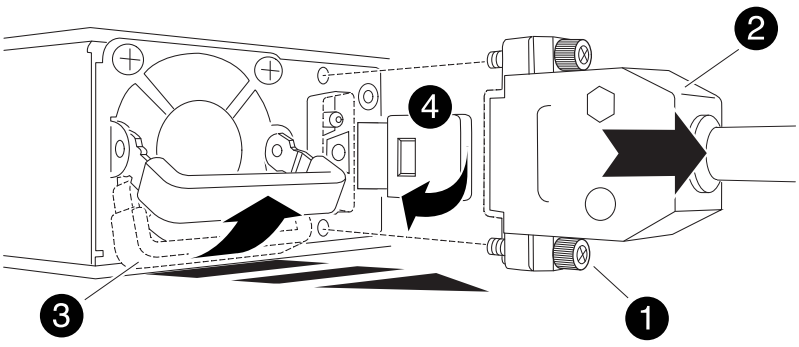

1	Thumb screws
2	D-SUB DC power cable connector

3	Power supply handle
4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF C250

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

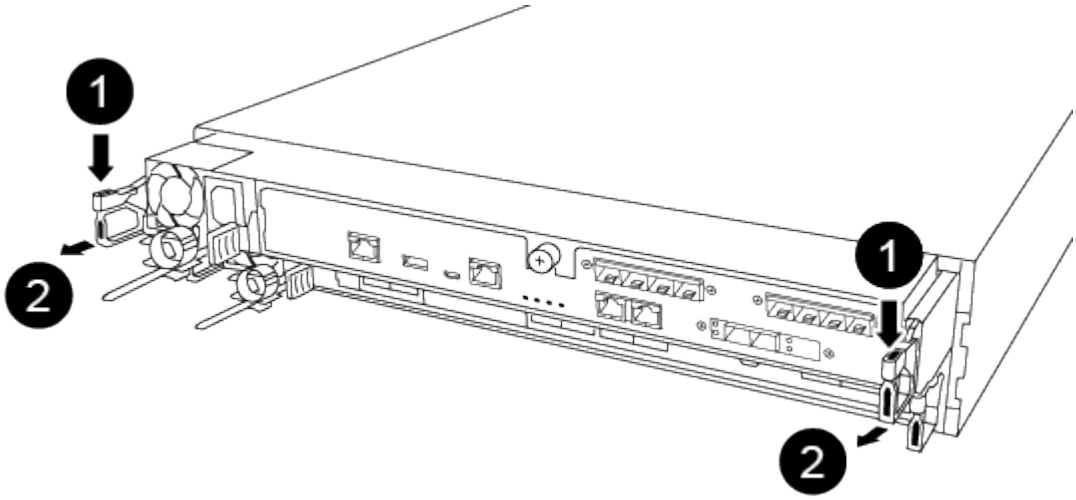

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

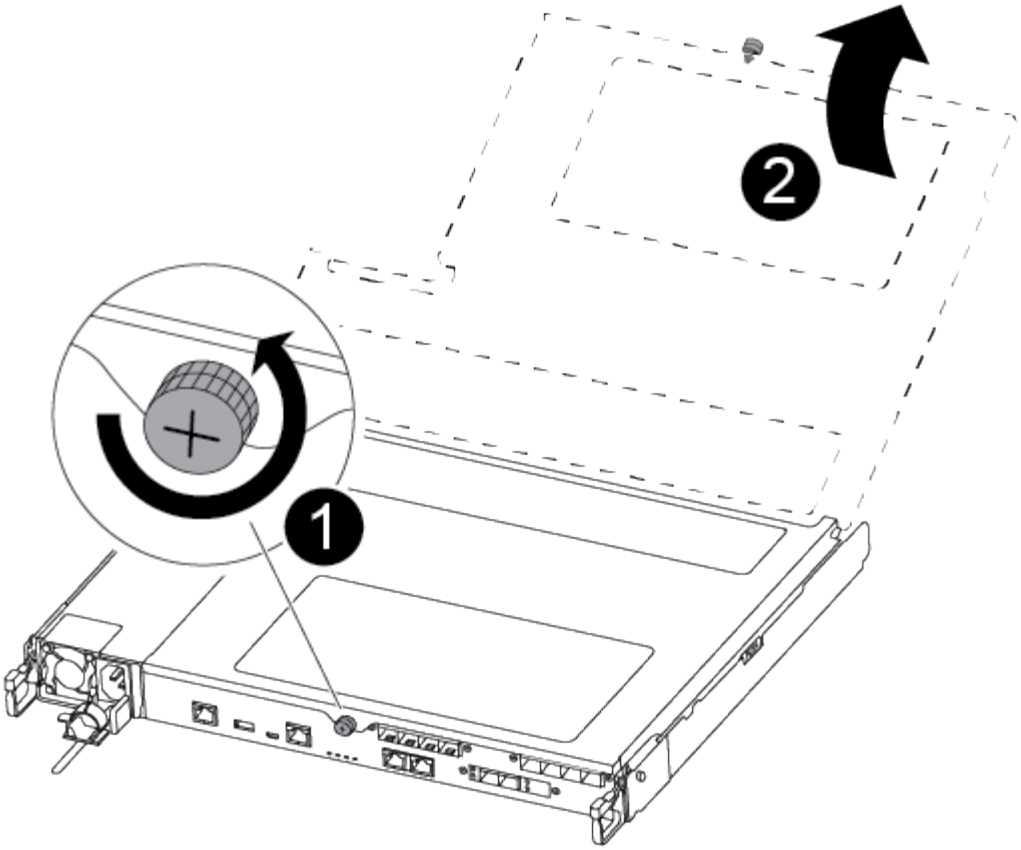

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

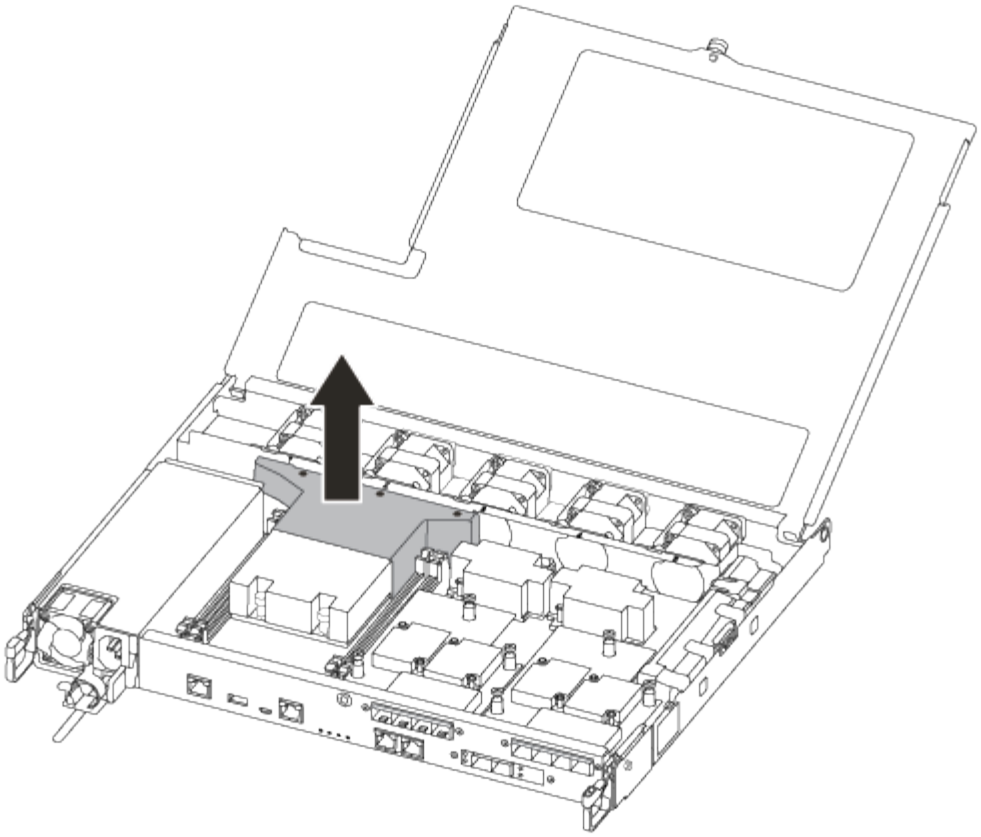

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

Use the following video or the tabulated steps to replace the RTC battery:

Animation - Replace the RTC battery

1. Locate the RTC battery between the heatsink and the midplane and remove it exactly as shown in the graphic.

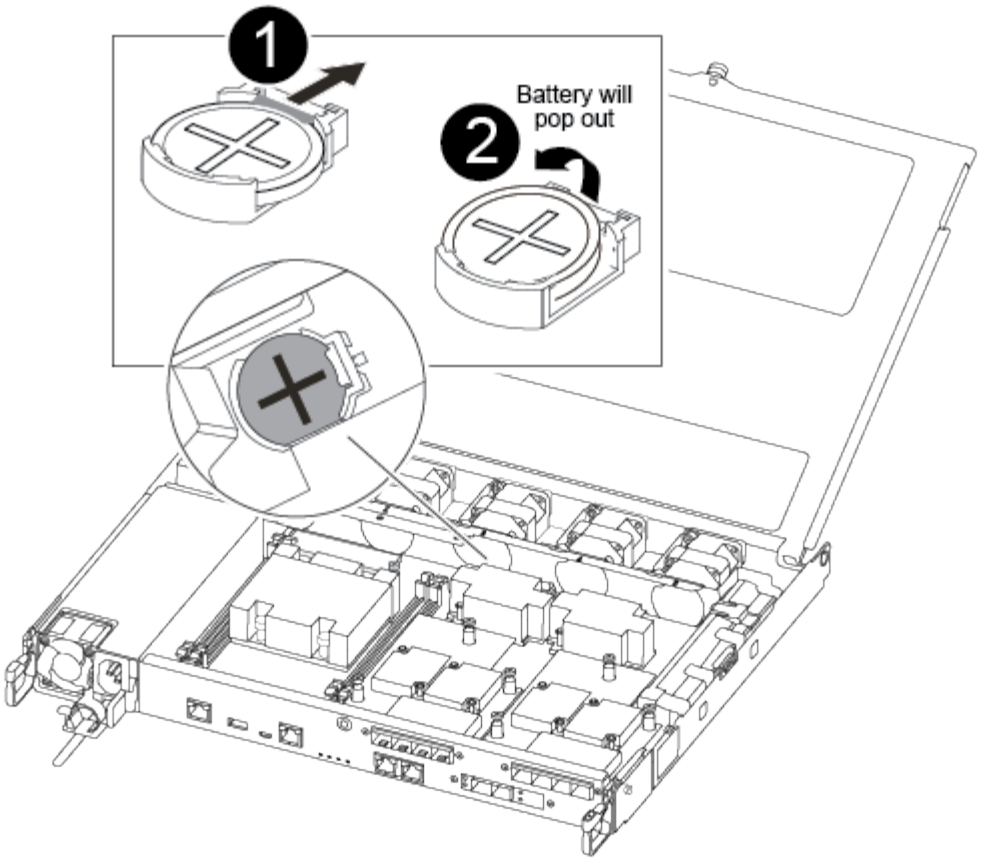

1	Gently pull tab away from the battery housing. **Attention:** Pulling it away aggressively might displace the tab.
2	Lift the battery up. **Note:** Make a note of the polarity of the battery.
3	The battery should eject out.

2. Remove the replacement battery from the antistatic shipping bag.
3. Locate the RTC battery holder between the heatsink and the midplane and insert it exactly as shown in the graphic.

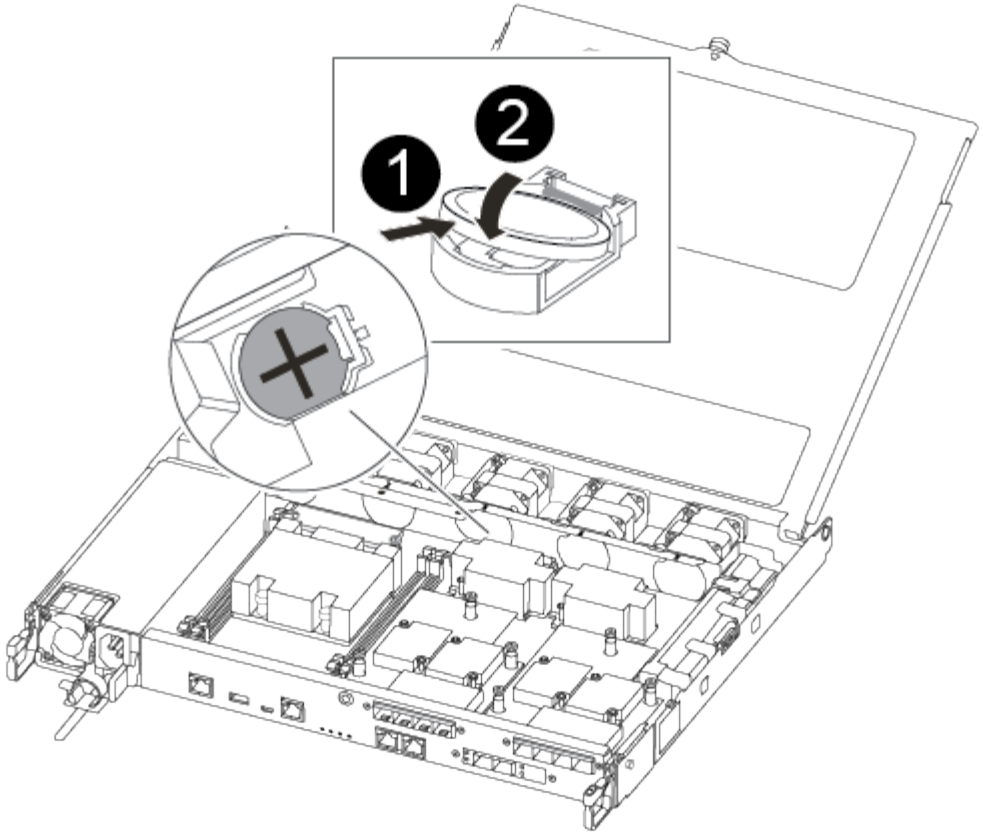

1	With positive polarity face up, slide the battery under the tab of the battery housing.
2	Push the battery gently into place and make sure the tab secures it to the housing. Pushing it in aggressively might cause the battery to eject out again.

4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Insert the controller module into the chassis:
    a. Ensure the latching mechanism arms are locked in the fully extended position.
    b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
    c. Place your index fingers through the finger holes from the inside of the latching mechanism.
    d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
    e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

       The controller module should be fully inserted and flush with the edges of the chassis.
    f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

       The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.
    g. Halt the controller at the LOADER prompt.
5. Reset the time and date on the controller:
    a. Check the date and time on the healthy controller with the `show date` command.
    b. At the LOADER prompt on the target controller, check the time and date.
    c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
    d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
    e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF C400 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

For MetroCluster configurations, see either:

- Install MetroCluster IP configuration
- Install MetroCluster Fabric-Attached configuration

Quick guide - AFF C400

The quick guide provides graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this procedure if you are familiar with installing NetApp systems.

Use the AFF C400 Installation and Setup Instructions.

The ASA C400 uses the same installation procedure as the AFF C400 system.

Video steps - AFF C400

The following video shows how to install and cable your new system.

Animation - AFF C400 Installation and setup instructions

If you have a MetroCluster configuration, use the MetroCluster installation content.

MetroCluster Documentation

Detailed guide - AFF C400

This guide gives detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

If you have a MetroCluster configuration, use the MetroCluster installation content.

MetroCluster Documentation

Step 1: Prepare for installation

To install your system, you need to create an account, register the system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

- You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release

Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

- You need to provide the following at your site:
  - Rack space for the storage system
  - Phillips #2 screwdriver
  - Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
100 GbE cable (QSFP28)	X66211A-05 (112-00595), 0.5m X66211A-1 (112-00573), 1m X66211A-2 (112-00574), 2m X66211A-5 (112-00574), 5m		Storage, cluster interconnect/HA, and Ethernet data (order-dependent)
25 GbE cable (SFP28)	X66240-2 (112-00598), 2m X66240-5 (112-00639), 5m		GbE network connection (order-dependent)
32 Gb FC (SFP+ Op)	X66250-2 (112-00342), 2m X66250-5 (112-00344), 5m X66250-15 (112-00346), 15m		FC network connection
Optical cables	X66250-2-N-C (112-00342)		16 Gb FC or 25GbE cables for mezzanine cards (order-dependent)

Type of cable…	Part number and length	Connector type	For…
RJ-45 (order dependent)	X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network
Micro-USB console cable	Not applicable		Console connection used during software setup if laptop or console does not support network discovery.
Power cables	Not applicable		Powering up the system

4. Review the *NetApp ONTAP Configuration Guide* and collect the required information listed in that guide.

   ONTAP Configuration Guide

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

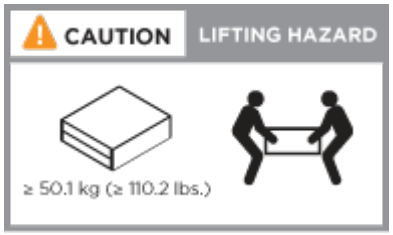

3. Attach cable management devices to the back of the controllers (as shown).

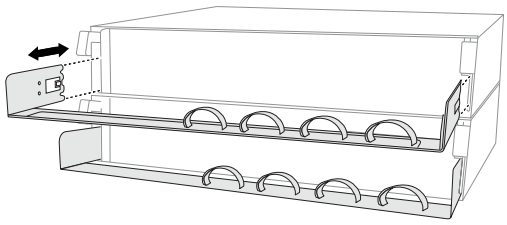

4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the switched cluster method.

About this task

- If the port labels on the card are not visible, you can identify the ports by checking the card installation orientation (for C400, the PCIe connector socket is on the left side of the card slot), and then look for the card by part number in NetApp Hardware Universe, which shows a graphic of the bezel with the port labels. You can find the card part number using the sysconfig -a command or on the system packing list.
- If you are cabling an MetroCluster IP configuration, ports e0a/e0b are available for hosting data LIFs (usually in Default IPSpace).

Option 1: Cable a two-node switchless cluster

A controller module's cluster interconnect and HA ports are cabled to its partner controller module. The optional data ports, optional NIC cards, and management ports on the controller modules are connected to switches.

Before you begin

You must have contacted your network administrator for information about connecting the system to the switches.

About this task

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.

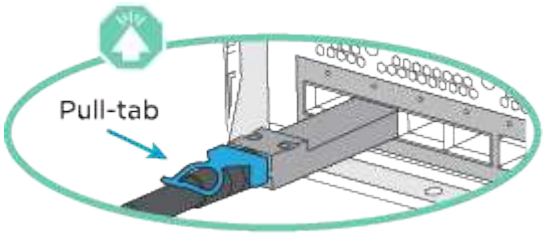

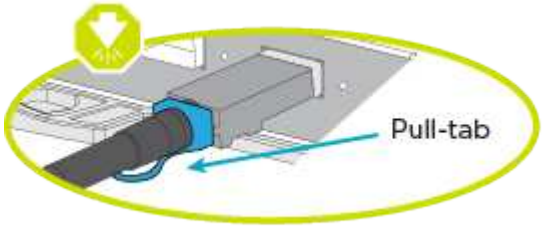

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the illustration to complete the cabling between the controllers and the switches:

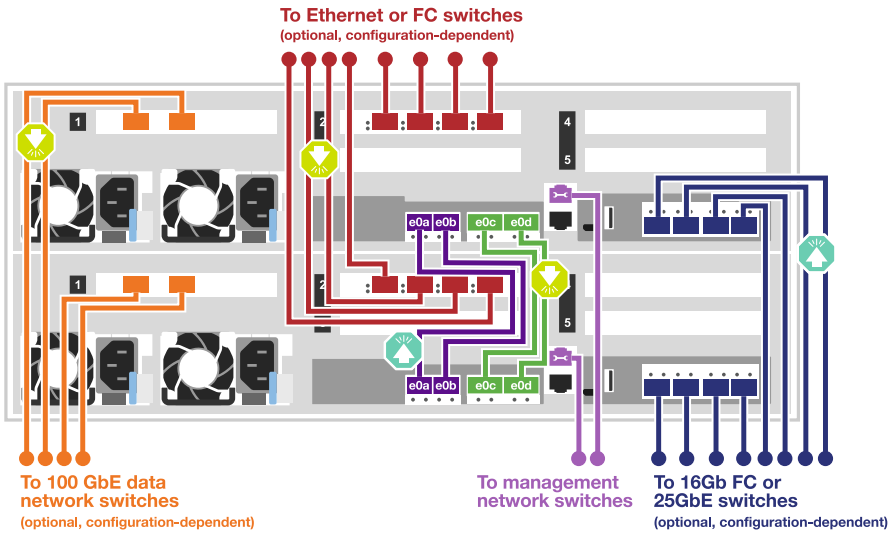

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Option 2: Cable a switched cluster

Controller module cluster interconnect and HA ports are cabled to the cluster/HA switch. The optional data ports, optional NIC cards, mezzanine cards, and management ports are connected to switches.

Before you begin

You must have contacted your network administrator for information about connecting the system to the switches.

About this task

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.

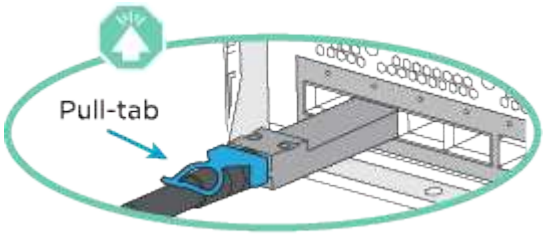

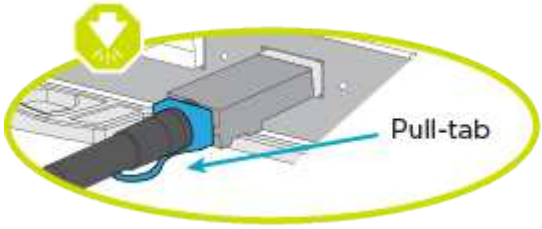

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the illustration to complete the cabling between the controllers and the switches:

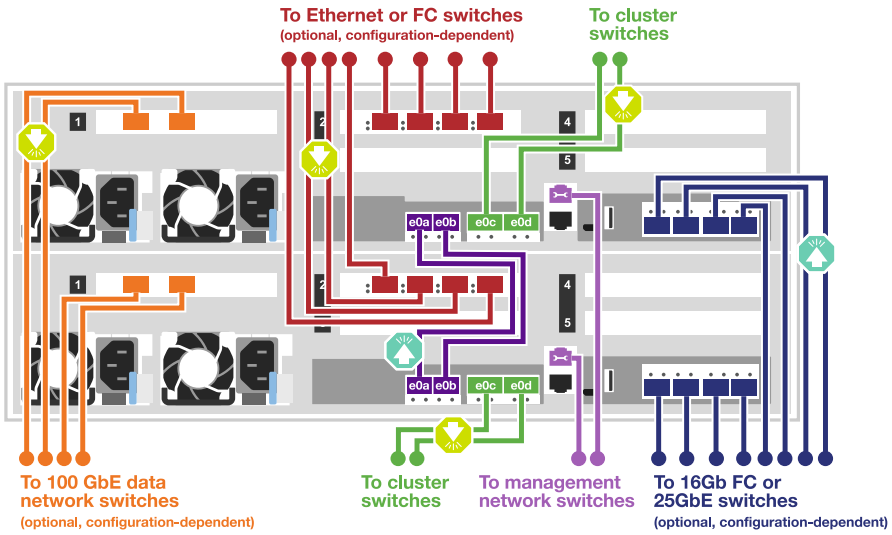

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Step 4: Cable controllers to drive shelves

The following options show you how to cable one or two NS224 drive shelves to your system.

Option 1: Cable the controllers to a single drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

About this task

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the NS224 are up.

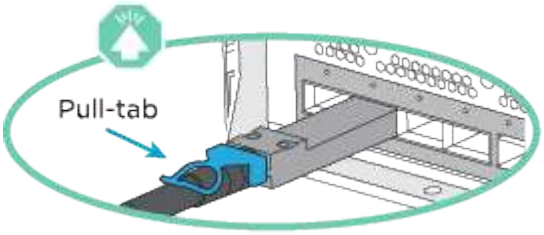

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following illustration to cable your controllers to a single drive shelf.

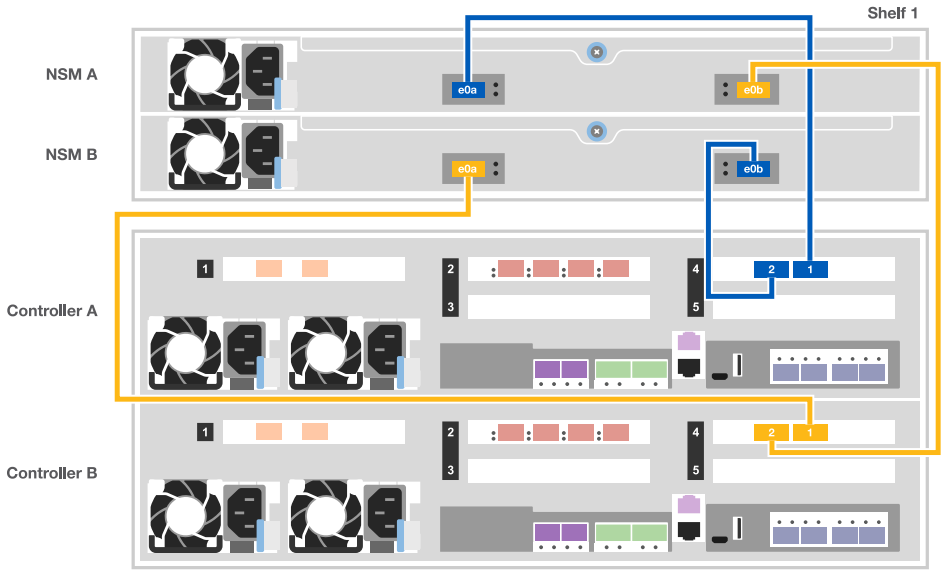

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Option 2: Cable the controllers to two drive shelves

You must cable each controller to the NSM modules on both NS224 drive shelves.

About this task

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the NS224 are up.

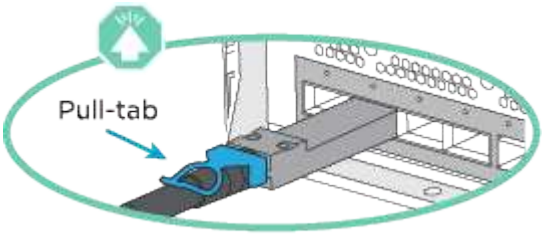

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following illustration to cable your controllers to two drive shelves.

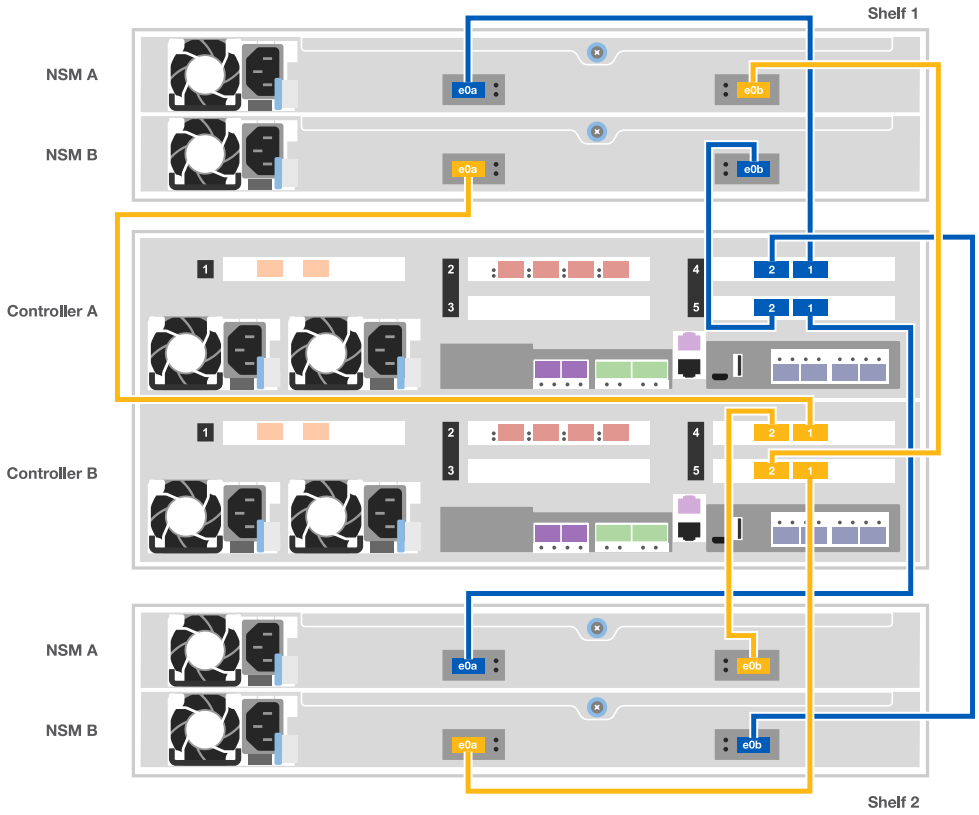

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Completing system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

1. Use the following animation to power on and set shelf IDs for one or more drive shelves:

   For NS224 drive shelves, shelf IDs are pre-set to 00 and 01. If you want to change the shelf IDs, use the straightened end of a paperclip, or narrow tipped ball point pen to access the shelf ID button behind the faceplate.

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Connect your laptop to the Management switch.

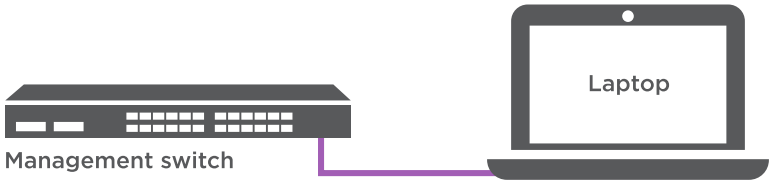

5. Select an ONTAP icon listed to discover:

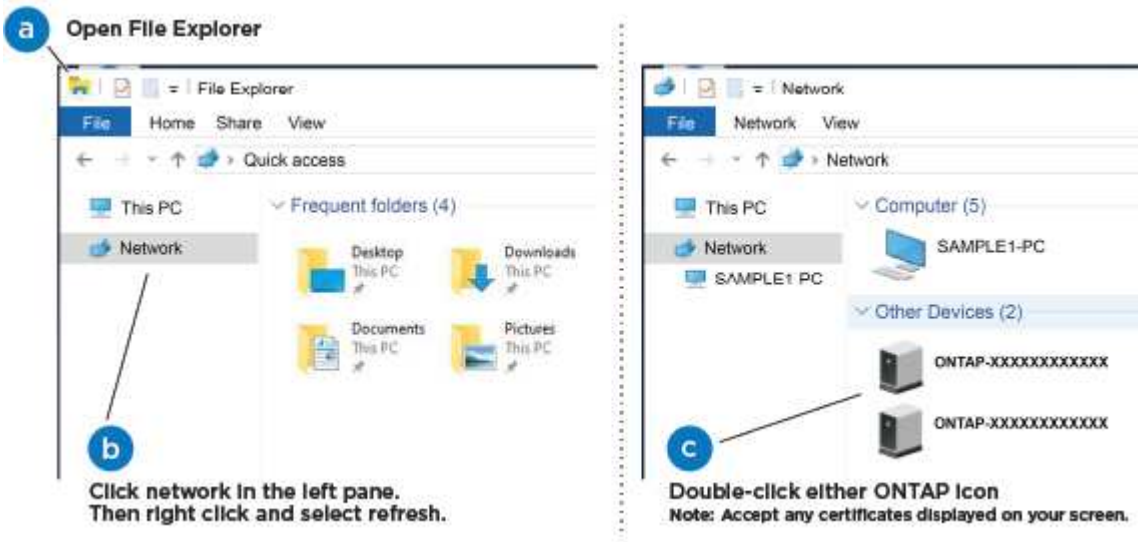

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

   XXXXX is the system serial number for the target node.

   System Manager opens.

6. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

ONTAP Configuration Guide

7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet .
   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Use the following animation to power on and set shelf IDs for one or more drive shelves:

   For NS224 drive shelves, shelf IDs are pre-set to 00 and 01. If you want to change the shelf IDs, use the straightened end of a paperclip, or narrow tipped ball point pen to access the shelf ID button behind the faceplate.

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   Initial booting may take up to eight minutes.

4. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.

If the management network has DHCP…	Then…
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the *NetApp ONTAP Configuration guide*.

      ONTAP Configuration Guide

6. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain AFF C400 hardware

Maintain the hardware of your AFF C400 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the AFF C400 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF C400 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.
NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.
PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - AFF C400

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on the partner node to reinstall ONTAP on the replacement boot media in your AFF C400 storage

system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1 **Review the boot media requirements**

Review the requirements for boot media replacement.

2 **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3 **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4 **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5 **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - AFF C400

Before replacing the boot media in your AFF C400, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:

- /cfcard/kmip/servers.cfg file.
- /cfcard/kmip/certs/client.crt file.
- /cfcard/kmip/certs/client.key file.
- /cfcard/kmip/certs/CA.pem file.

- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - AFF C400

Shut down the impaired controller in your AFF C400 storage system to prevent data loss and ensure system stability when replacing the boot media.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - AFF C400

The boot media in your AFF C400 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

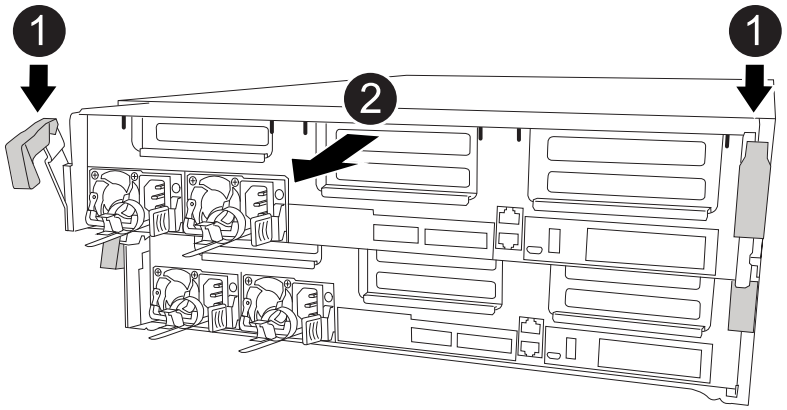

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. Open the air duct:

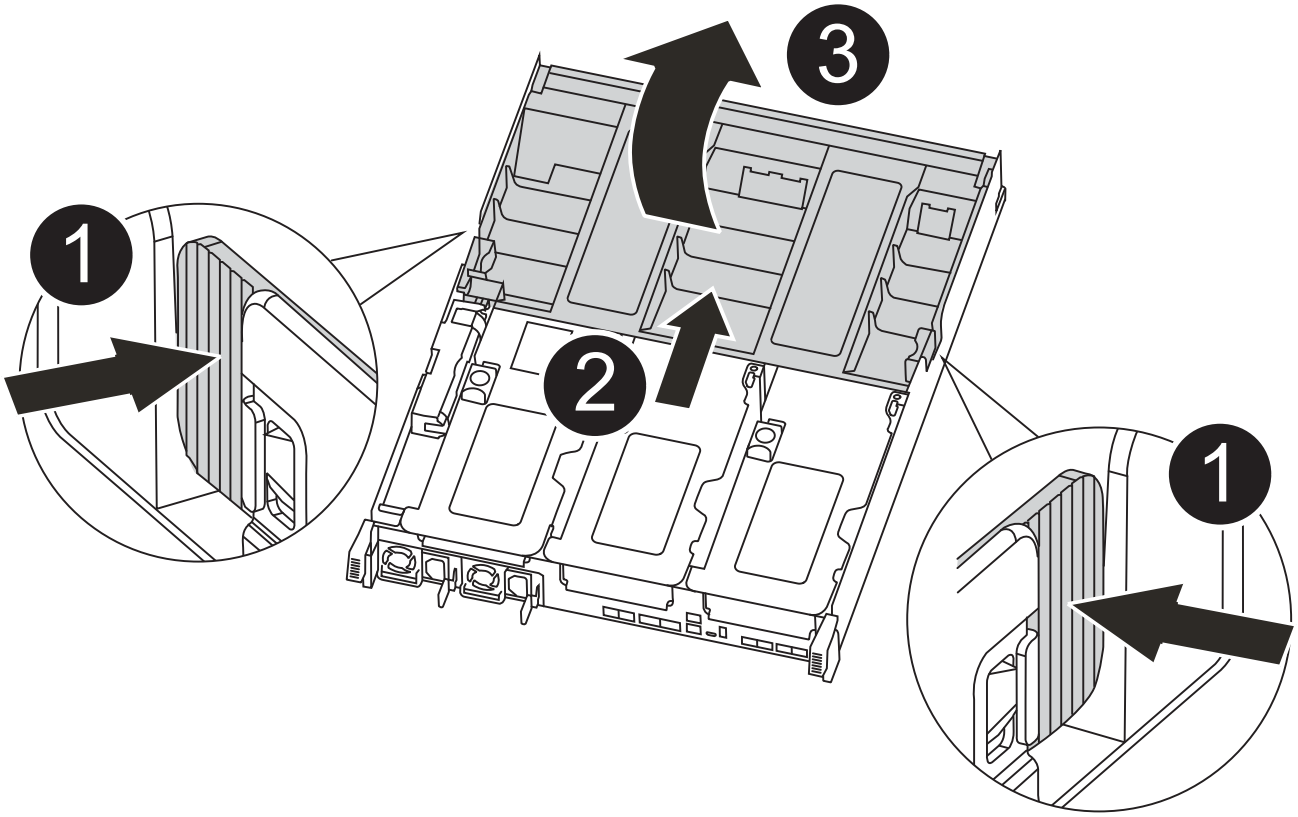

1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.

b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

9. Locate and remove the boot media from the controller module:

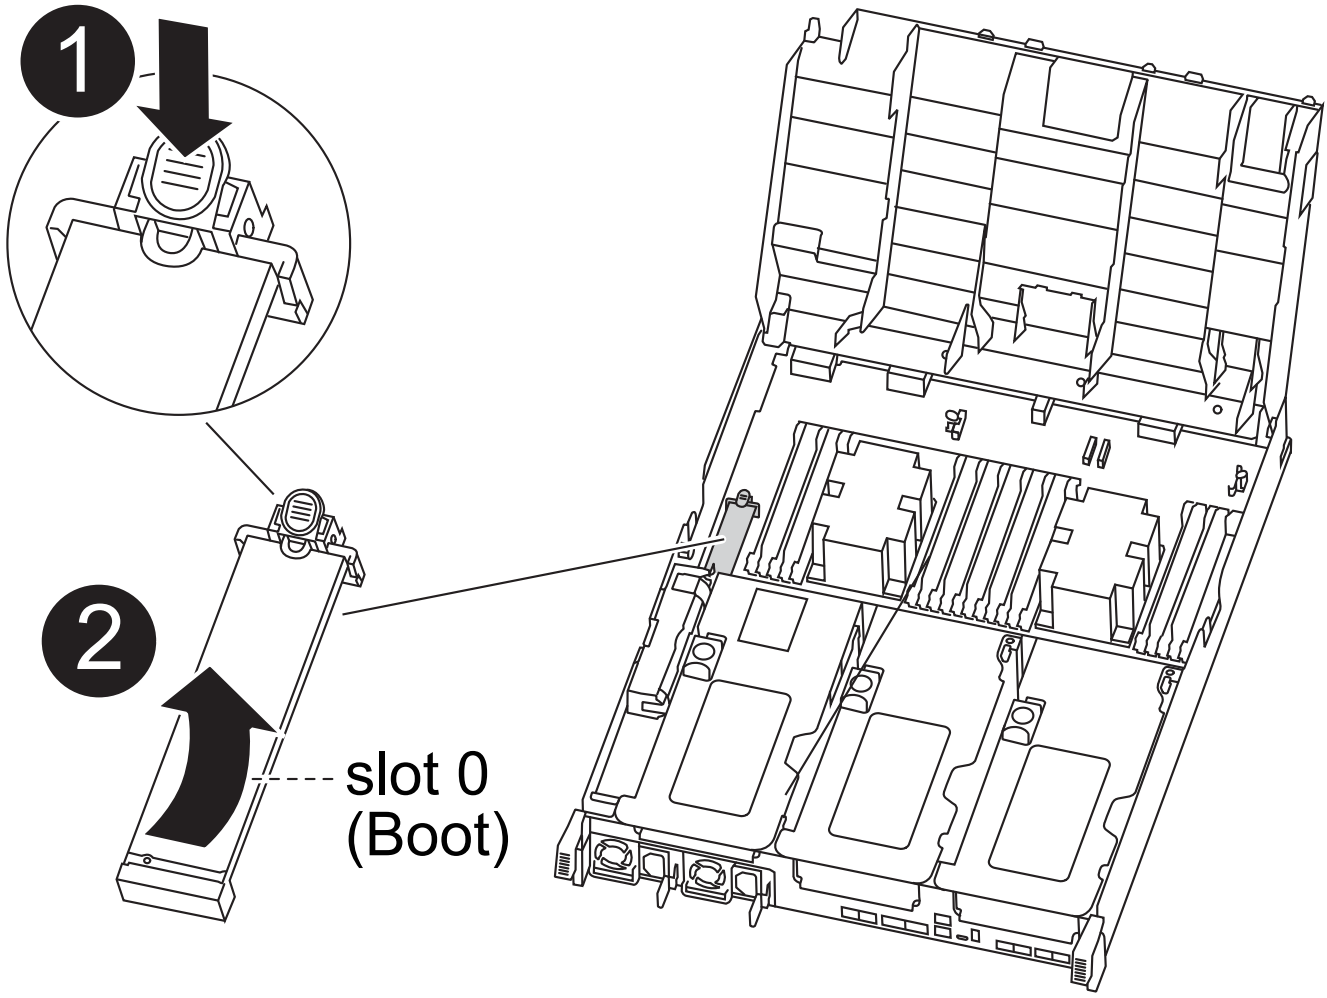

1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.

b. Rotate the boot media up and gently pull the boot media out of the socket.

10. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
11. Check the boot media to make sure that it is seated squarely and completely in the socket.

    If necessary, remove the boot media and reseat it into the socket.

12. Lock the boot media in place:
    a. Rotate the boot media down toward the motherboard.
    b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
    c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
13. Close the air duct.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - AFF C400

After installing the new boot media device in your AFF C400 system, you can start the automated boot media recovery process to restore the configuration from the partner node.

During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed part to NetApp - AFF C400

If a component in your AFF C400 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF C400

Get started with replacing the boot media in your AFF C400 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - AFF C400

Before replacing the boot media in your AFF C400 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - AFF C400

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - AFF C400

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster configuration

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 3: Controller is in a two-node Metrocluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes

that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the boot media and prepare for manual boot recovery - AFF C400

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

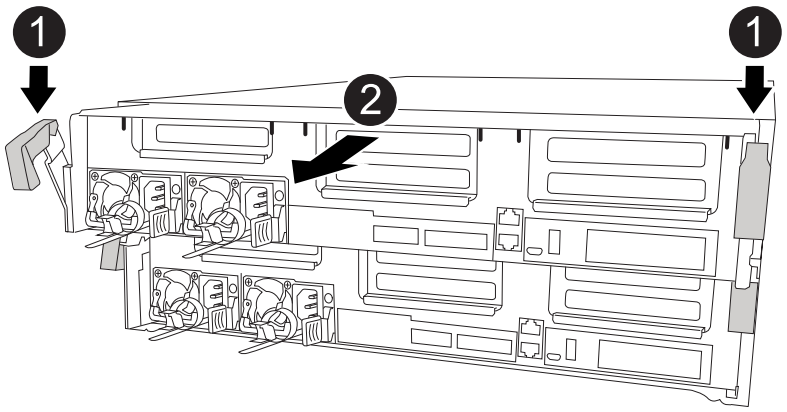

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 2: Replace the boot media

You must locate the boot media in the controller module (see the FRU map on the controller module), and then follow the directions to replace it.

Before you begin

Although the contents of the boot media is encrypted, it is a best practice to erase the contents of the boot media before replacing it. For more information, see the Statement of Volatility for your system on the NetApp Support Site.

> You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the boot media.

Animation - Replace the boot media

Steps

1. Open the air duct:

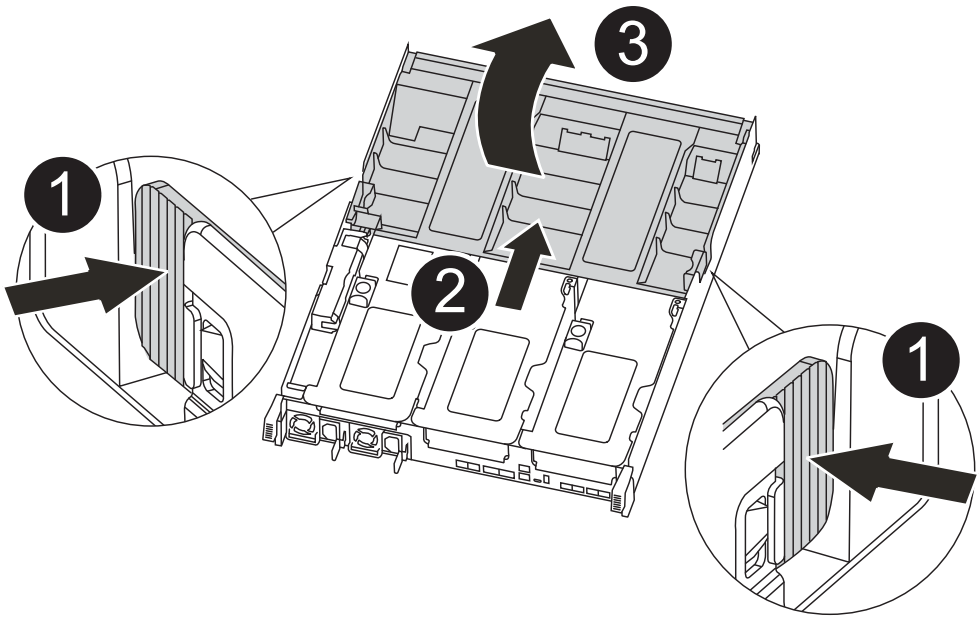

1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

2. Locate and remove the boot media from the controller module:

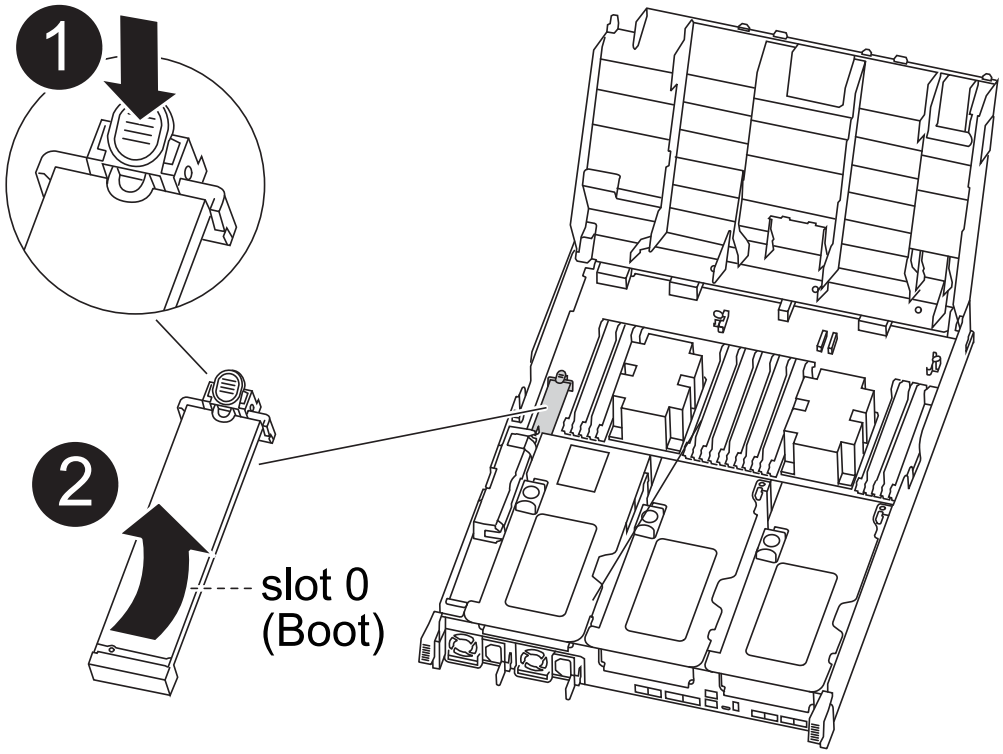

1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
b. Rotate the boot media up and gently pull the boot media out of the socket.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Lock the boot media in place:
   a. Rotate the boot media down toward the motherboard.
   b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
   c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
6. Close the air duct.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed does not have a boot image, so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the `var` file system.

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      There are two folders in the unzipped service image file:

      - `boot`
      - `efi`

   c. Copy the `efi` folder to the top directory on the USB flash drive.

      If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

      The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

   d. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.
5. Plug the power cable into the power supply and reinstall the power cable retainer.
6. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.
7. Complete the installation of the controller module:

a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

   The locking latches rise when the controller module is fully seated.

   > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

d. If you have not already done so, reinstall the cable management device.

8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then `halt` the controller to boot to LOADER.

9. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:

   a. Boot to Maintenance mode: `boot_ontap maint`

   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t` *`initiator adapter_name`*

   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - AFF C400

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.

3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `...would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore encryption - AFF C400

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
********************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
********************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - AFF C400

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF C400

To replace the chassis, you must move the fans and controller modules from the impaired chassis to the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial

outage in a multinode cluster.

Shut down the controllers - AFF C400

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Shut down the controllers when replacing a chassis

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Shut down a controller in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the

`-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Move the fans, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move the fans

To move the fan modules to the replacement chassis when replacing the chassis, you must perform a specific sequence of tasks.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.
4. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.
5. Set the fan module aside.
6. Repeat the preceding steps for any remaining fan modules.

7. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.
10. Repeat these steps for the remaining fan modules.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
   b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - AFF C400

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for *HA-state* can be one of the following:

      - `ha`
      - `mcc`
      - `mcc-2n`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

4. Reinstall the bezel on the front of the system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller module

Overview of controller module replacement - AFF C400

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement node* is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - AFF C400

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the controller module hardware - AFF C400

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

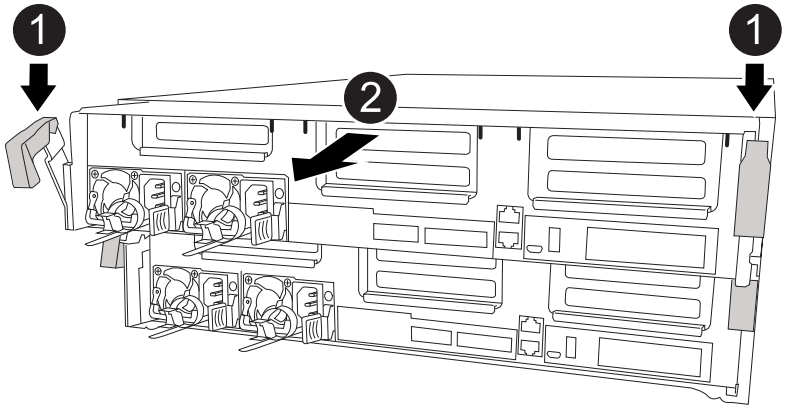

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. On the replacement controller module, open the air duct and remove the empty risers from the controller module using the animation, illustration, or the written steps:

Animation - Remove the empty risers from the replacement controller module

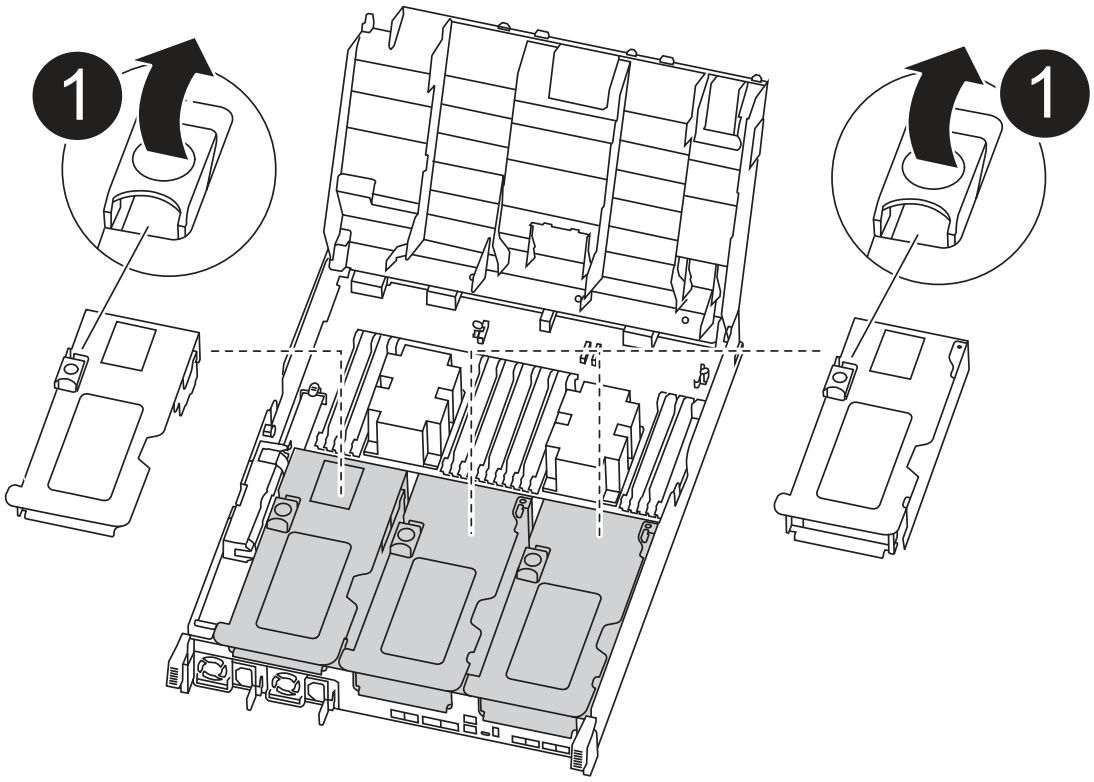

1	Riser latches

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
c. Rotate the riser locking latch on the left side of riser 1 up and toward air duct, lift the riser up, and then set it aside.
d. Repeat the previous step for the remaining risers.

Step 2: Move the power supplies

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

You can use the following animation, illustration, or the written steps to move the power supplies to the replacement controller module.

Animation - Move the power supplies

1. Remove the power supply:

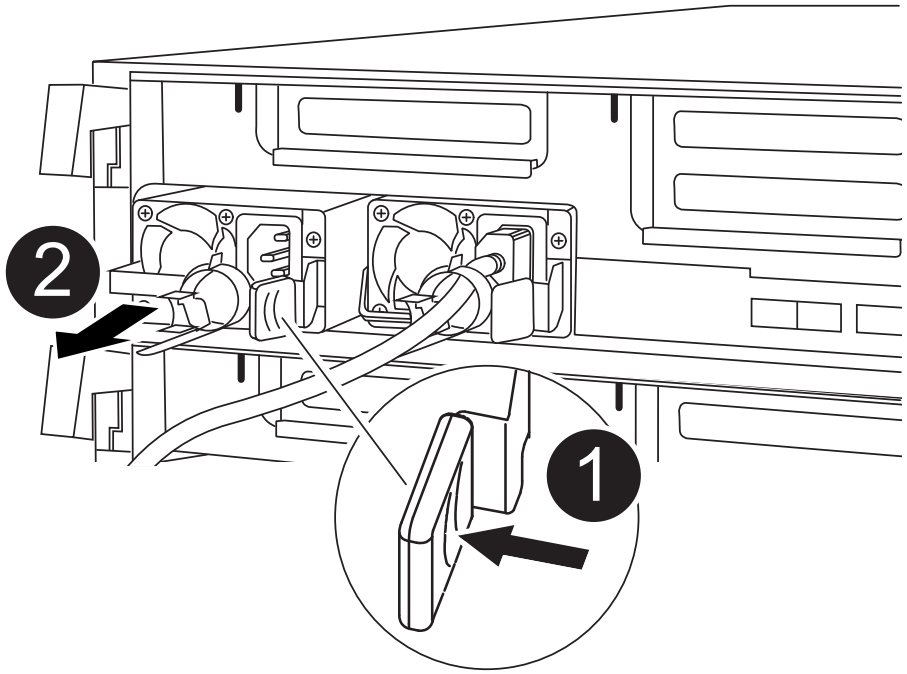

1	PSU locking tab
2	Power cable retainer

a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.

b. Press the blue locking tab to release the power supply from the chassis.

c. Using both hands, pull the power supply out of the chassis, and then set it aside.

1. Move the power supply to the new controller module, and then install it.
2. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

3. Repeat the preceding steps for any remaining power supplies.

Step 3: Move the NVDIMM battery

To move the NVDIMM battery from the impaired controller module to the replacement controller module, you must perform a specific sequence of steps.

You can use the following animation, illustration, or the written steps to move the NVDIMM battery from the impaired controller module to the replacement controller module.

Animation - Move the NVDIMM battery

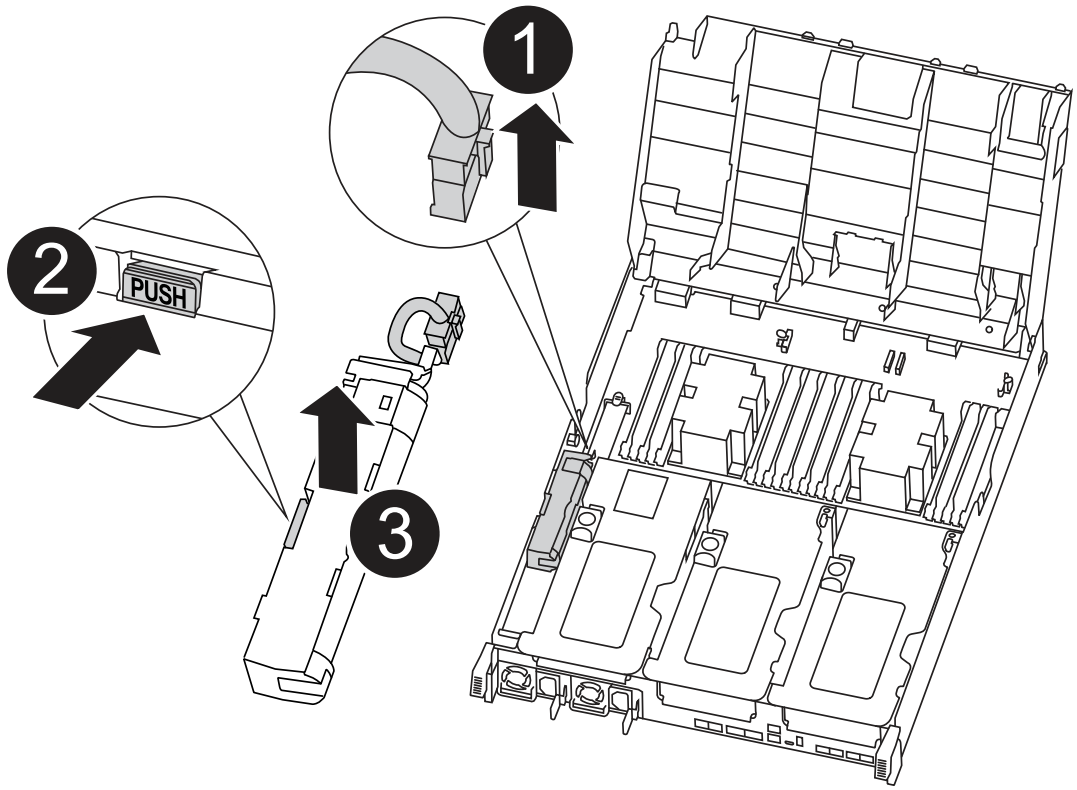

1	NVDIMM battery plug
2	NVDIMM battery locking tab
3	NVDIMM battery

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.
3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.

Do not plug the battery cable back into the motherboard until instructed to do so.

Step 4: Move the boot media

You must locate the boot media, and then follow the directions to remove it from the impaired controller module and insert it into the replacement controller module.

You can use the following animation, illustration, or the written steps to move the boot media from the impaired controller module to the replacement controller module.

Animation - Move the boot media

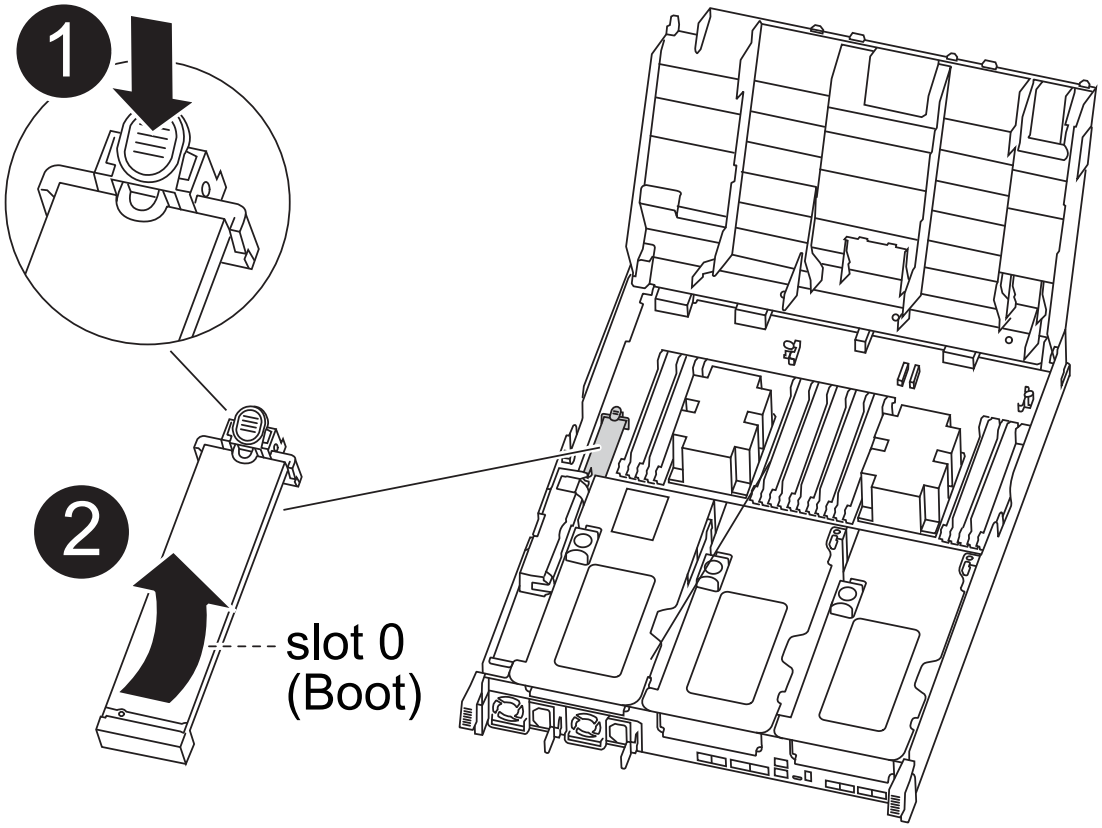

1	Boot media locking tab
2	Boot media

1. Locate and remove the boot media from the controller module:
   a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
   b. Rotate the boot media up and gently pull the boot media out of the socket.
2. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
3. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

4. Lock the boot media in place:

a. Rotate the boot media down toward the motherboard.

b. Press the blue locking button so that it is in the open position.

c. Placing your fingers at the end of the boot media by the blue button, firmly push down on the boot media end to engage the blue locking button.

Step 5: Move the PCIe risers and mezzanine card

As part of the controller replacement process, you must move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

You can use the following animations, illustrations, or the written steps to move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

Moving PCIe riser 1 and 2 (left and middle risers):

Animation - Move PCI risers 1 and 2

Moving the mezzanine card and riser 3 (right riser):

Animation - Move the mezzanine card and riser 3

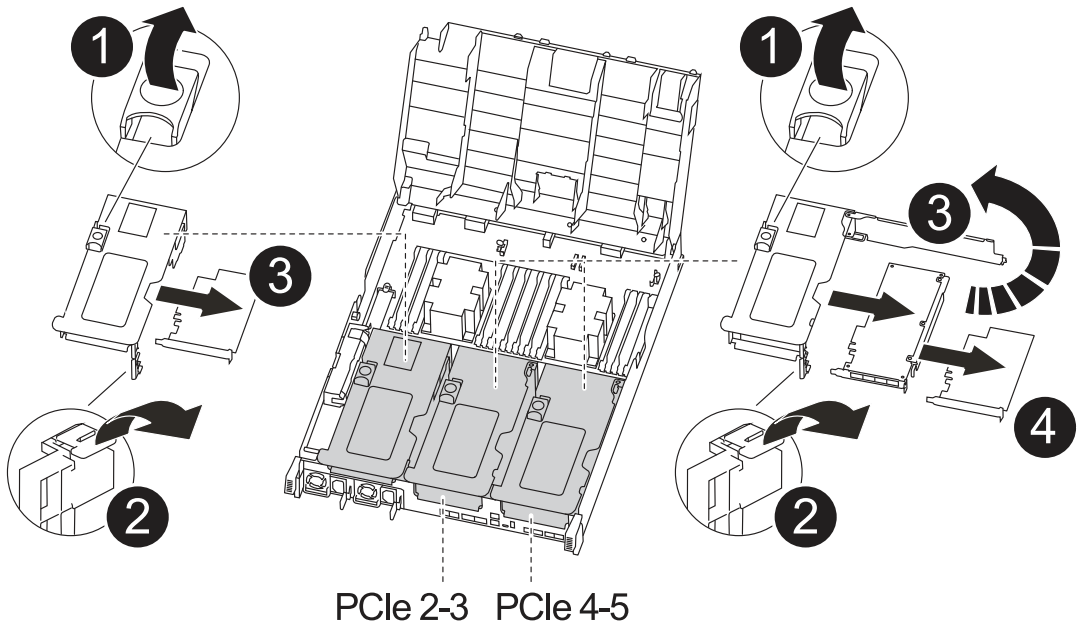

1	Riser locking latch
2	PCI card locking latch
3	PCI locking plate
4	PCI card

1. Move PCIe risers one and two from the impaired controller module to the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then move it to the replacement controller module.
   d. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins, push the riser squarely into the socket on the motherboard, and then rotate the latch down flush with the sheet metal on the riser.
   e. Repeat this step for riser number 2.
2. Remove riser number 3, remove the mezzanine card, and install both into the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then set it aside on a stable, flat surface.
   d. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket, and then move it to the replacement controller module.
   e. Install the mezzanine in the replacement controller and secure it with the thumbscrews.
   f. Install the third riser in the replacement controller module.

Step 6: Move the DIMMs

You need to locate the DIMMs, and then move them from the impaired controller module to the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

You can use the following animation, illustration, or the written steps to move the DIMMs from the impaired controller module to the replacement controller module.

Animation - Move the DIMMs

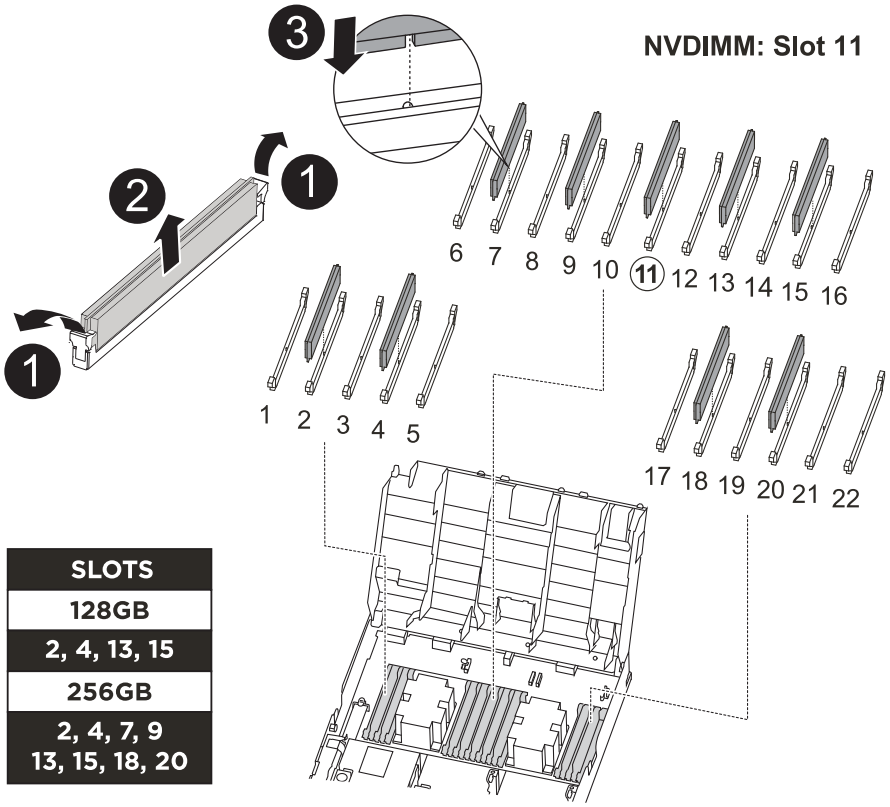

1	DIMM locking tabs
2	DIMM
3	DIMM socket

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Verify that the NVDIMM battery is not plugged into the new controller module.
4. Move the DIMMs from the impaired controller module to the replacement controller module:

   Make sure that you install the each DIMM into the same slot it occupied in the impaired controller module.

   a. Eject the DIMM from its slot by slowly pushing apart the DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

      Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   b. Locate the corresponding DIMM slot on the replacement controller module.

   c. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the

DIMM squarely into the socket.

The DIMMs fit tightly in the socket, but should go in easily. If not, realign the DIMM with the socket and reinsert it.

d. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

e. Repeat these substeps for the remaining DIMMs.

5. Plug the NVDIMM battery into the motherboard.

Make sure that the plug locks down onto the controller module.

Step 7: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

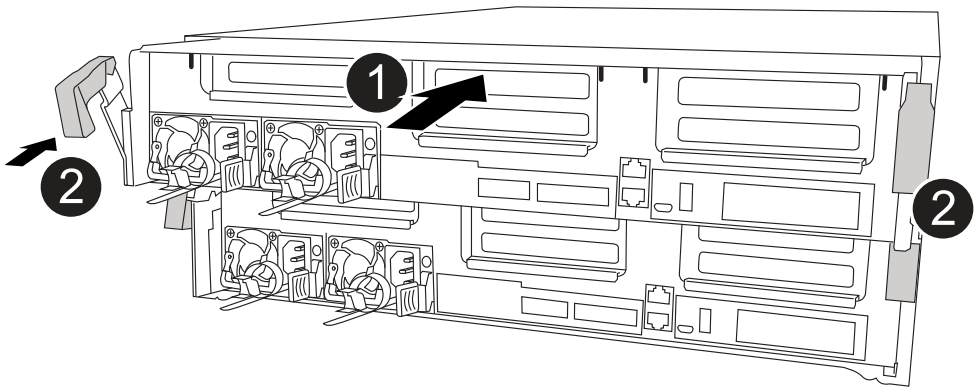

1	Slide controller into the chassis
2	Locking latches

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:

a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

Restore and verify the system configuration - AFF C400

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF C400

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.

c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.

d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`

   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *`local-node-name`* `partner savecore -s`

d. Return to the admin privilege level: `set -privilege admin`

5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at

which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The *replacement* controller is the current owner of the disks on the disaster site.

     Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF C400

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no

configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    ------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF C400

You must replace a DIMM in the controller module when your system registers an increasing number of correctable error correction codes (ECC); failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

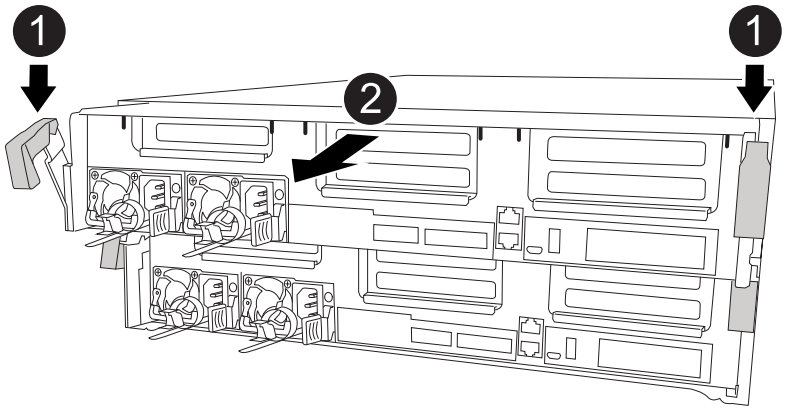

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace system DIMMs

Replacing a system DIMM involves identifying the target DIMM through the associated error message, locating the target DIMM using the FRU map on the air duct, and then replacing the DIMM.

You can use the following animation, illustration, or the written steps to replace a system DIMM.

> The animation and illustration show empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace a system DIMM

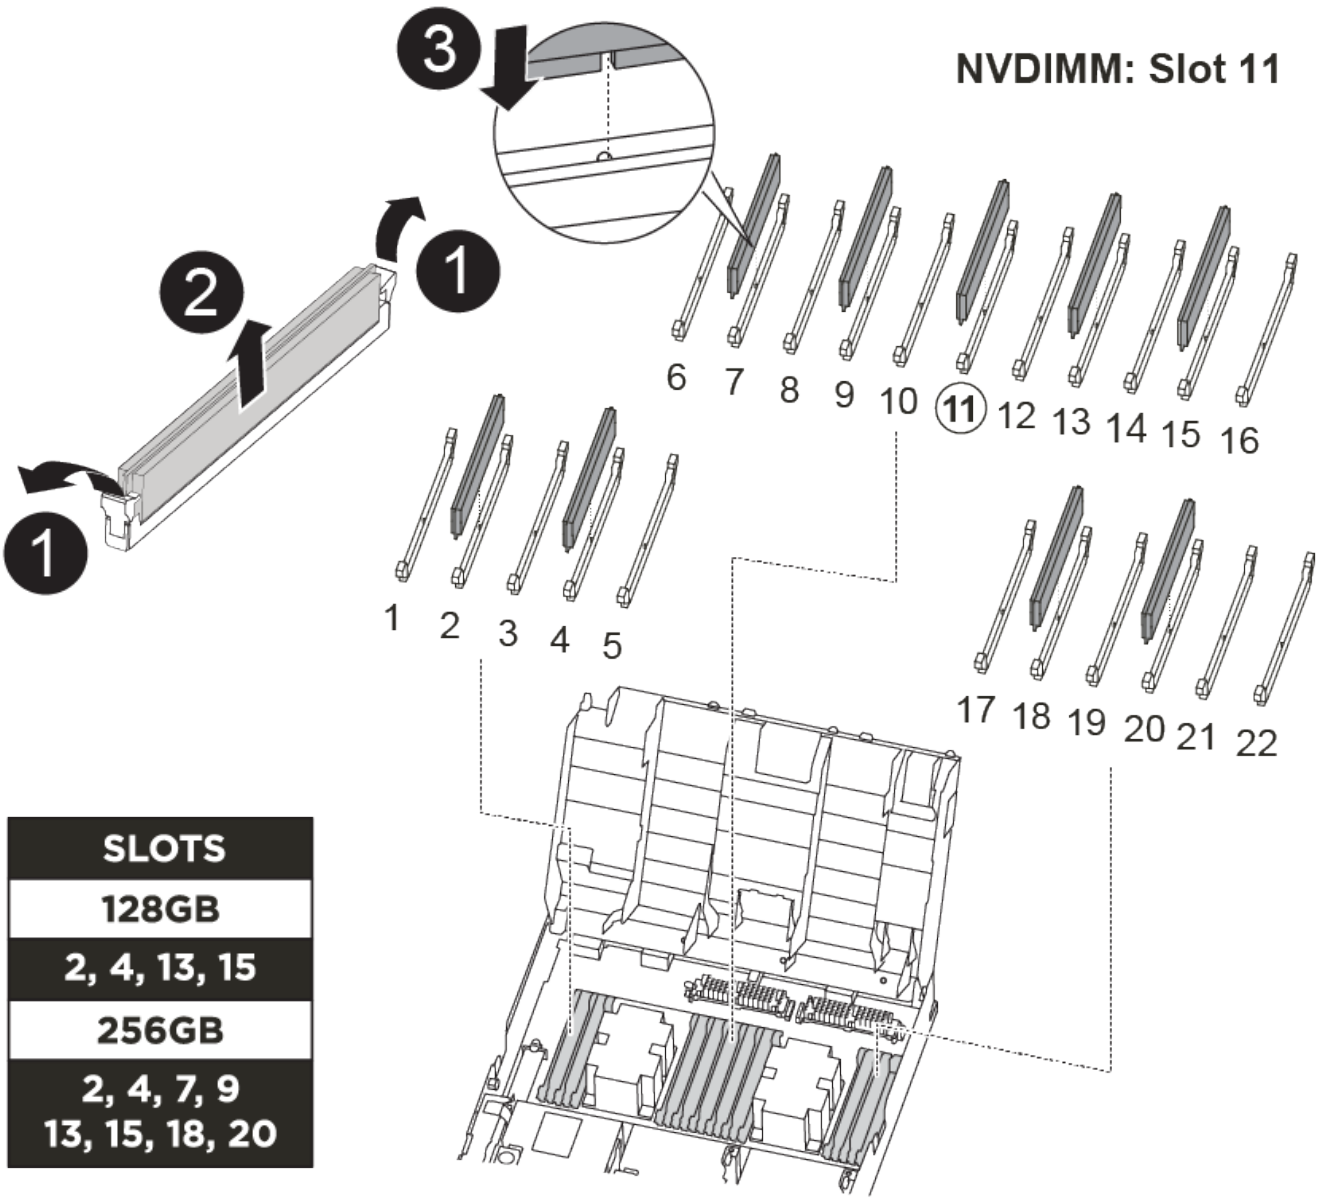

1	DIMM locking tabs
2	DIMM
3	DIMM socket

The DIMMs are located in sockets 2, 4, 13, and 15. The NVDIMM is located in slot 11.

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely

open position.

2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
4. Eject the DIMM from its socket by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the socket.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

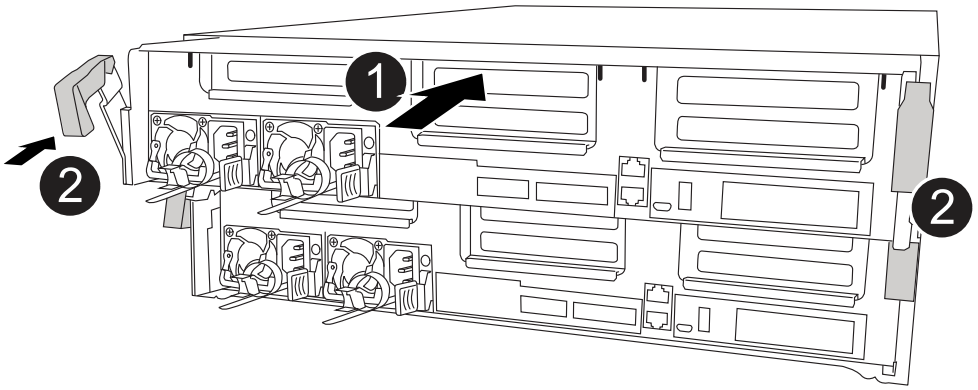

1	Controller module
2	Controller locking latches

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
   b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.
   g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto`

```
-giveback true
```

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    ------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a fan module - AFF C400

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

You can use the following animation, illustration, or the written steps to hot-swap a fan module.

Animation - Replace a fan

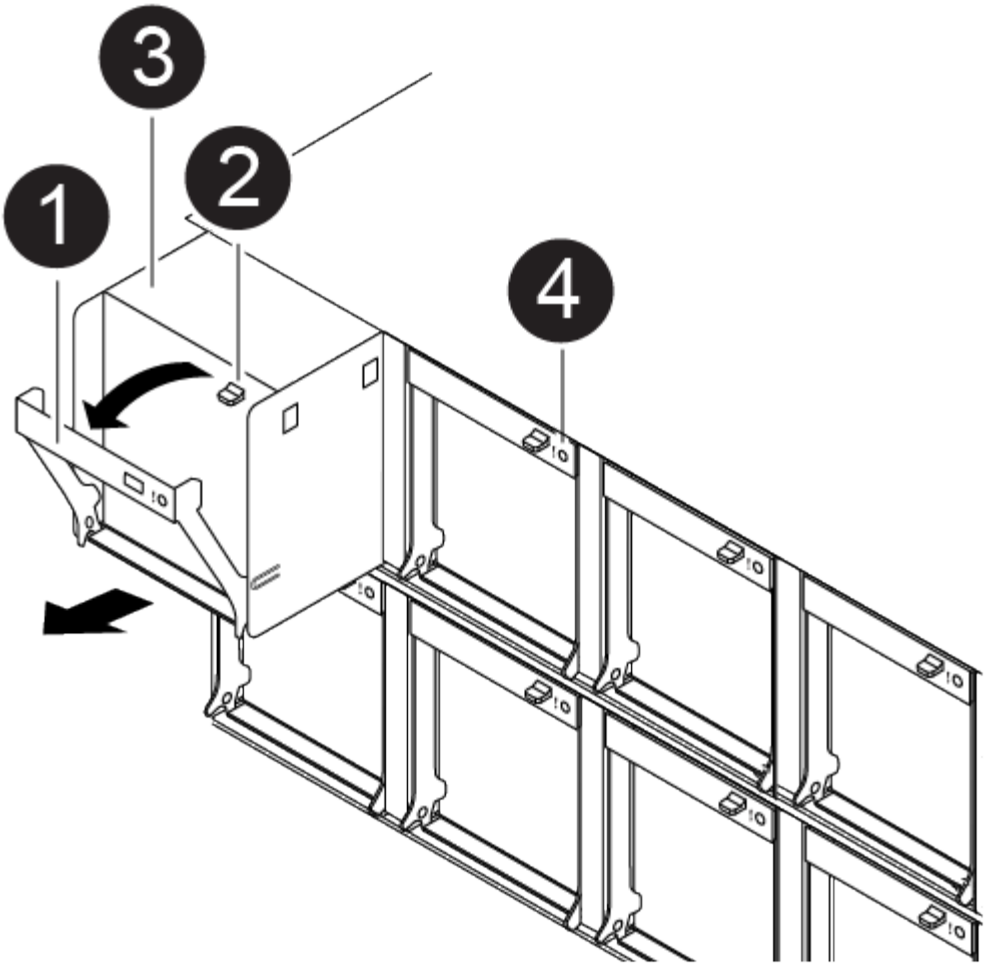

1	Fan handle
2	Locking tab
3	Fan
4	Status LED

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.

5. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

6. Set the fan module aside.
7. Insert the replacement fan module into the chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The Attention LED should not be lit after the fan is seated and has spun up to operational speed.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - AFF C400

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

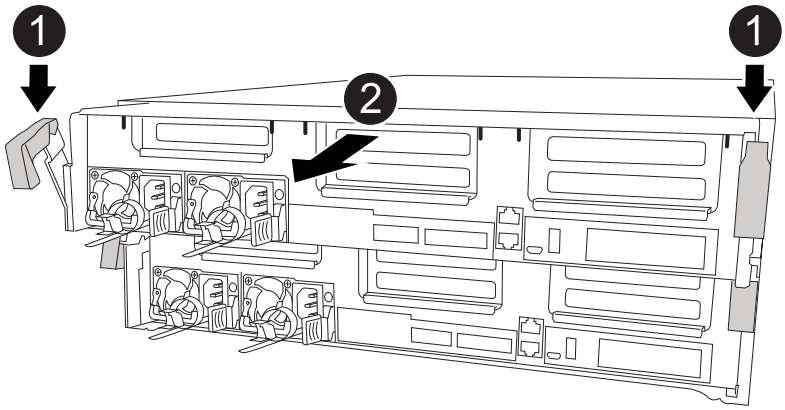

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module. See the FRU map inside the controller module to locate the NVDIMM battery.

The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.

You can use the following animation, illustration, or the written steps to replace the NVDIMM battery.

Animation - Replace the NVDIMM battery

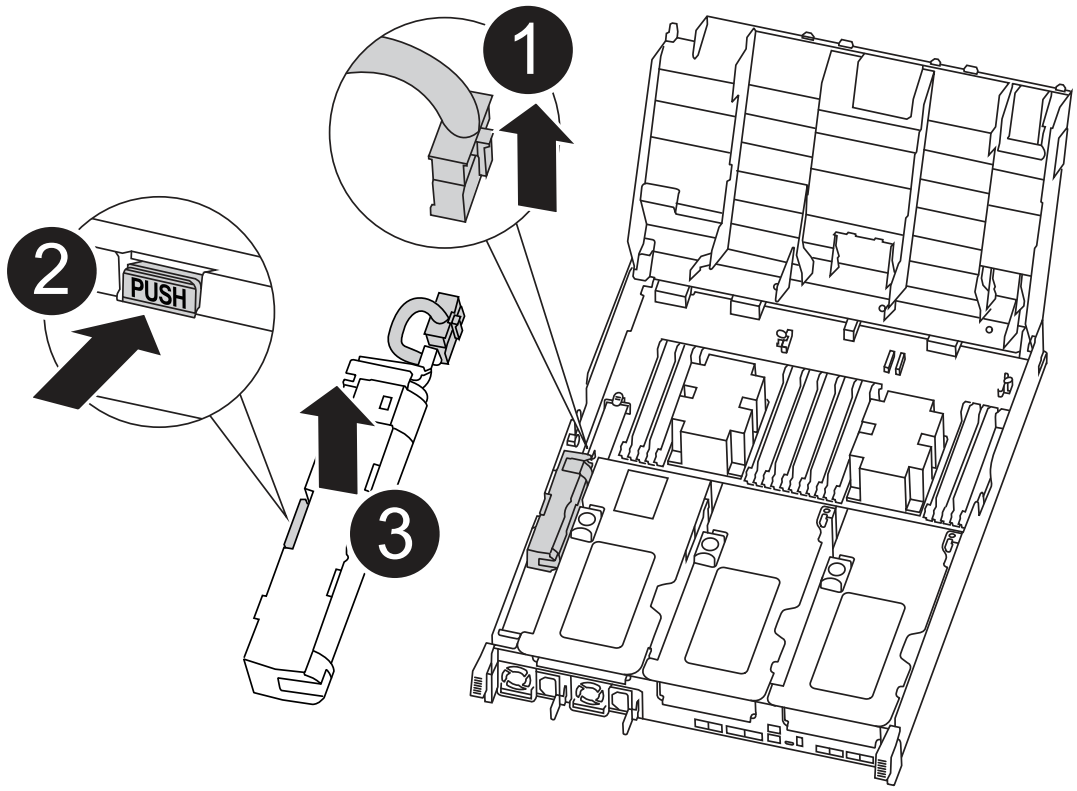

1	Battery plug
2	Locking tab
3	NVDIMM battery

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.
3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Remove the replacement battery from its package.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.
7. Plug the battery plug back into the controller module, and then close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

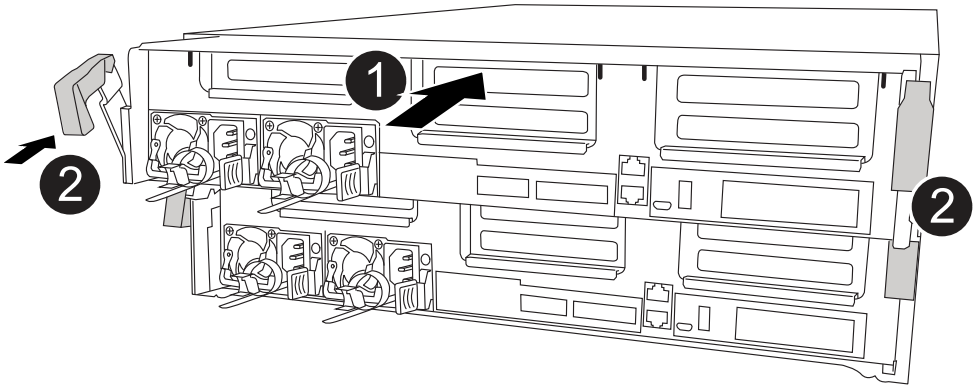

1	Controller module
2	Controller locking latches

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
   b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to

interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an NVDIMM - AFF C400

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

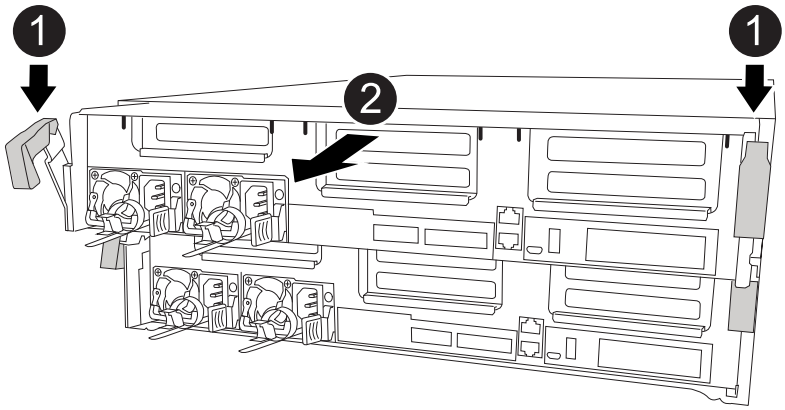

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the FRU map on top of the air duct or the FRU Map on the top of the slot 1 riser.

- The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.
- Although the contents of the NVDIMM is encrypted, it is a best practice to erase the contents of the NVDIMM before replacing it. For more information, see the Statement of Volatility on the NetApp Support

Site.

You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the NVDIMM.

The animation shows empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace the NVDIMM

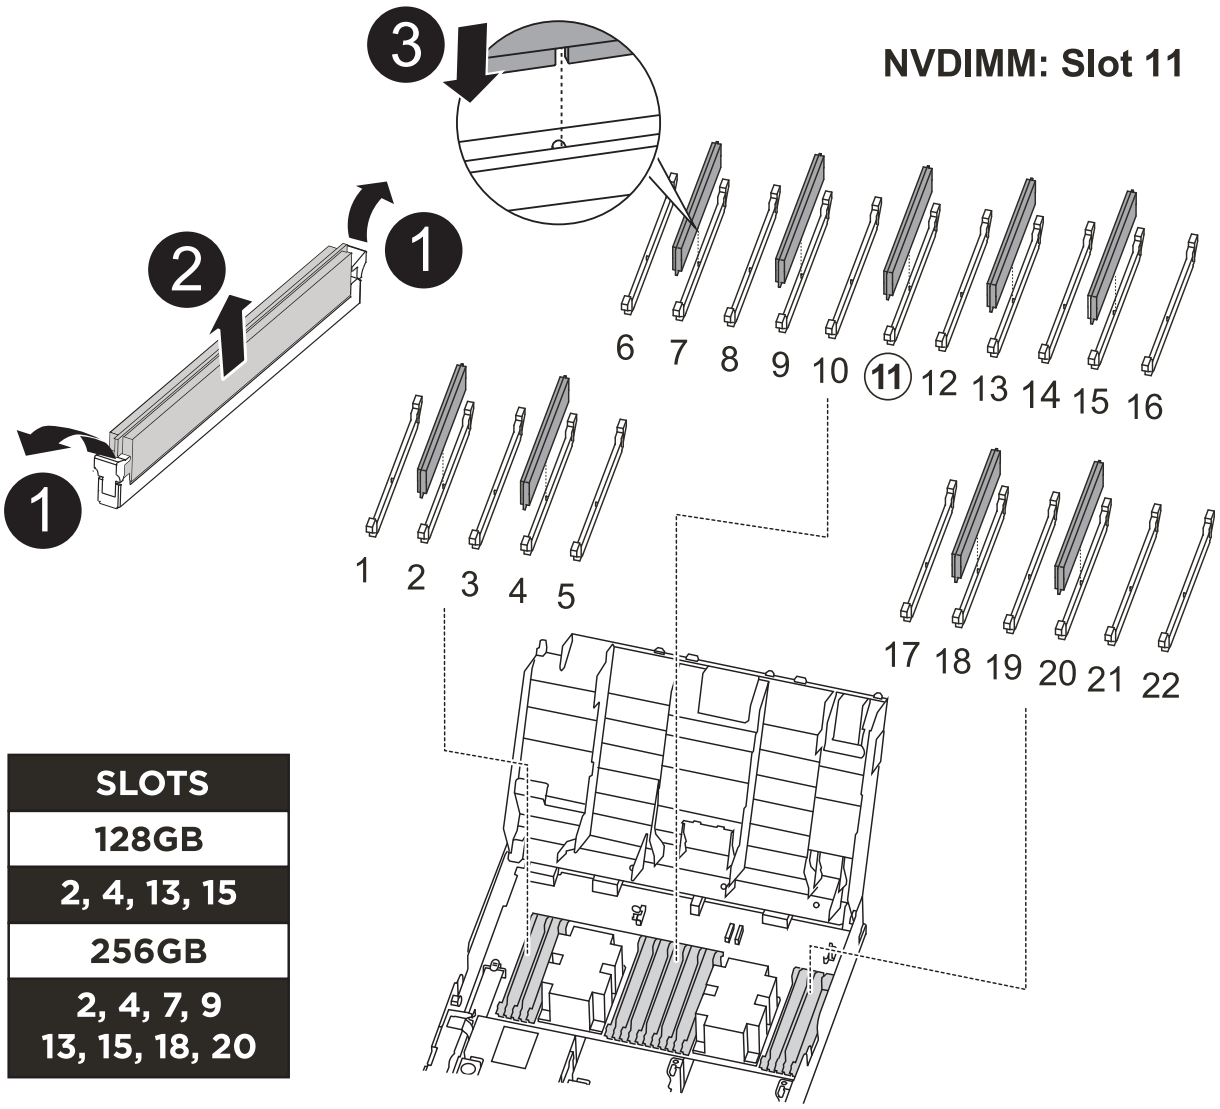

1	DIMM locking tabs
2	DIMM

3	DIMM socket

1. Open the air duct and then locate the NVDIMM in slot 11 on your controller module.

   The NVDIMM looks significantly different than system DIMMs.

2. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

3. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

4. Locate the slot where you are installing the NVDIMM.
5. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
7. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe or mezzanine card - AFF C400

To replace a PCIe or mezzanine card, you must disconnect the cables and any SFP and QSFP modules from the cards, replace the failed PCIe or mezzanine card, and then recable the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft

vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

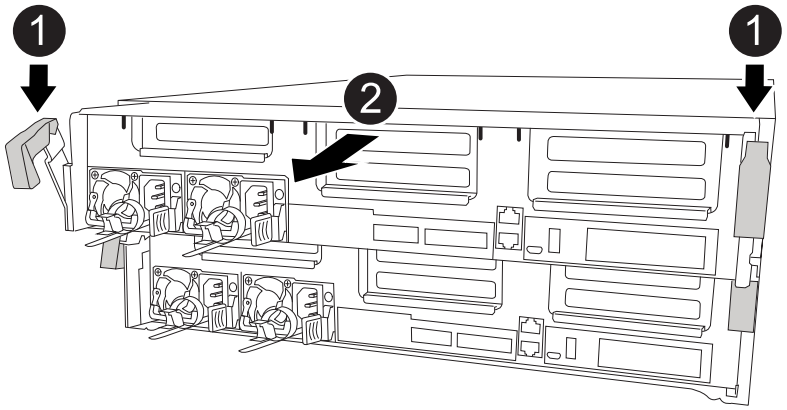

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace a PCIe card

To replace a PCIe card, you must locate the failed PCIe card, remove the riser that contains the card from the controller module, replace the card, and then reinstall the PCIe riser in the controller module.

You can use the following animation, illustration, or the written steps to replace a PCIe card.

Animation - Replace a PCIe card

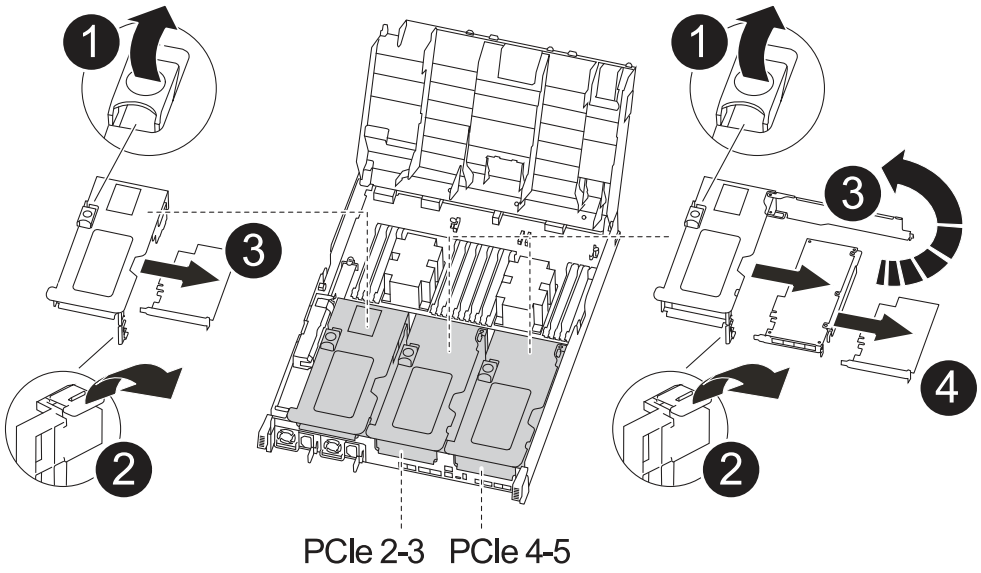

1	Riser locking latch
2	PCI card locking latch
3	PCI locking plate
4	PCI card

1. Remove the riser containing the card to be replaced:
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up straight up and set it aside on a stable flat surface,
2. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. For risers 2 and 3 only, swing the side panel up.
   d. Remove the PCIe card from the riser by gently pushing up on the bracket and lift the card straight out of the socket.
3. Install the replacement PCIe card in the riser by aligning the card with the socket, press the card into the socket and then close the side panel on the riser, if present.

Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The PCIe card must be fully and evenly seated in the slot.

> If you are installing a card in the bottom slot and cannot see the card socket well, remove the top card so that you can see the card socket, install the card, and then reinstall the card you removed from the top slot.

4. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 4: Replace the mezzanine card

The mezzanine card is located under riser number 3 (slots 4 and 5). You must remove that riser to access the mezzanine card, replace the mezzanine card, and then reinstall riser number 3. See the FRU map on the controller module for more information.

You can use the following animation, illustration, or the written steps to replace the mezzanine card.

Animation - Replace the mezzanine card

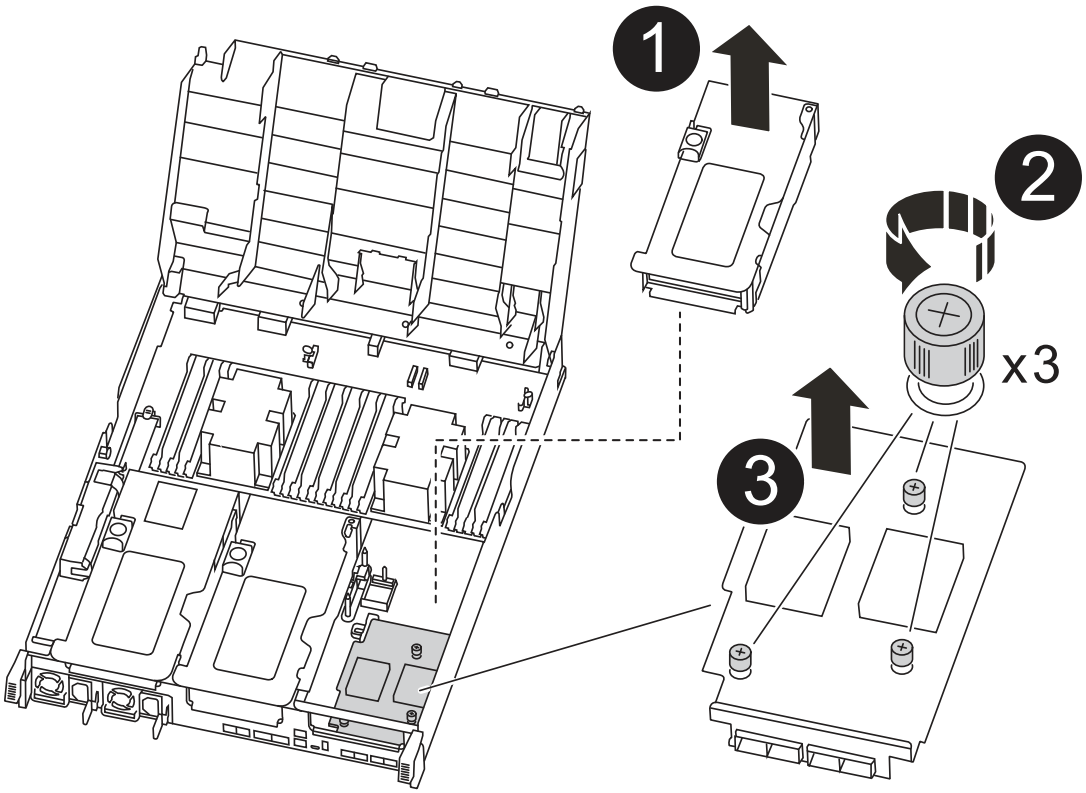

1	PCI riser

2	Riser thumbscrew
3	Riser card

1. Remove riser number 3 (slots 4 and 5):
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up, and then set it aside on a stable, flat surface.
2. Replace the mezzanine card:
   a. Remove any QSFP or SFP modules from the card.
   b. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket and set it aside.
   c. Align the replacement mezzanine card over the socket and the guide pins and gently push the card into the socket.
   d. Tighten the thumbscrews on the mezzanine card.
3. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 5: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

   The locking latches rise when the controller module is fully seated.

   > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

   > If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

5. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
6. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Restore the controller module to operation

To restore the controller, you must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 7: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 8: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replacing a power supply - AFF C400

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting the replacement PSU to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

You can use the following illustration with the written steps to replace the power supply.

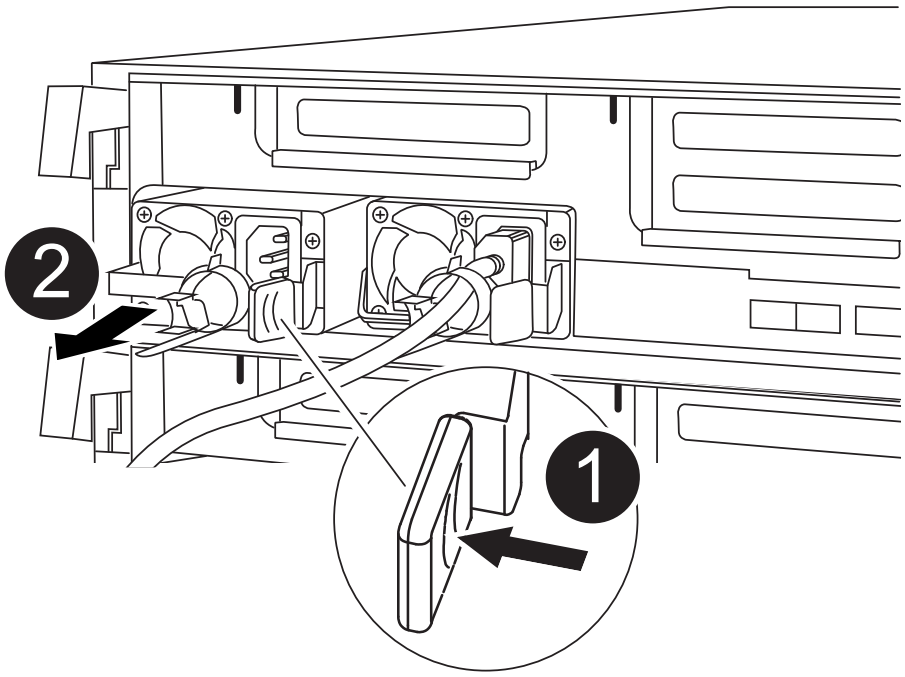

1	PSU locking tab
2	Power cable retainer

1. If you are not already grounded, properly ground yourself.

2. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
3. Disconnect the power supply:
   a. Open the power cable retainer, and then unplug the power cable from the power supply.
   b. Unplug the power cable from the power source.
4. Remove the power supply:
   a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
   b. Press the blue locking tab to release the power supply from the chassis.
   c. Using both hands, pull the power supply out of the chassis, and then set it aside.
5. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

6. Rotate the cam handle so that it is flush against the power supply.
7. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF C400

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

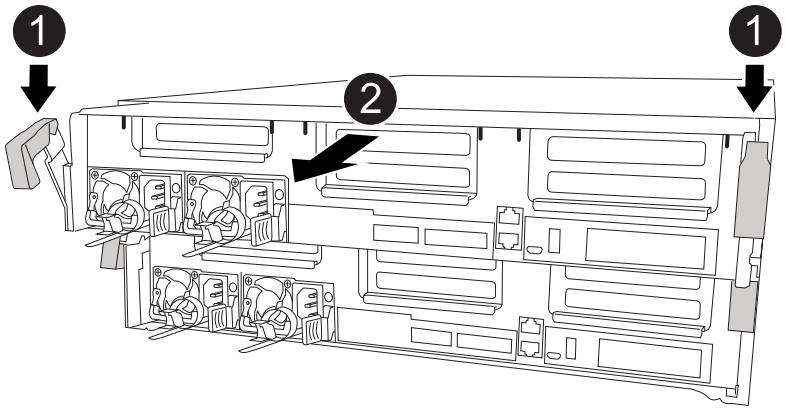

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the RTC battery

You need to locate the RTC battery inside the controller module, and then follow the specific sequence of steps. See the FRU map inside the controller module for the location of the RTC battery.

You can use the following animation, illustration, or the written steps to replace the RTC battery.

Animation- Replace the RTC battery

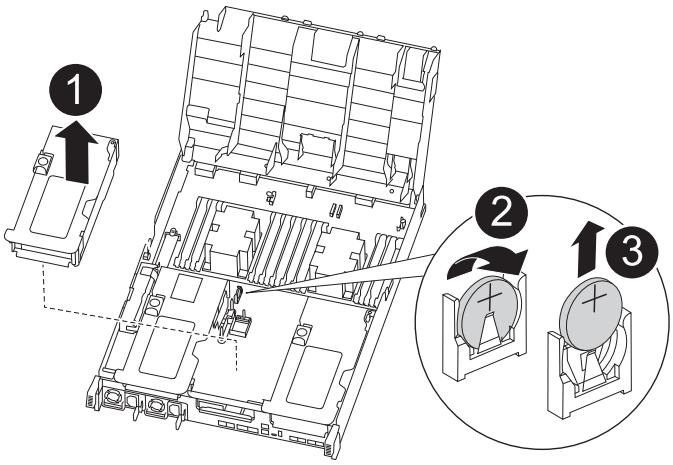

1	Middle riser
2	Remove RTC battery
3	Seat RTC battery

1. If you are not already grounded, properly ground yourself.
2. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
3. Locate, remove, and then replace the RTC battery:
   a. Using the FRU map, locate the RTC battery on the controller module.
   b. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

   > Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

   c. Remove the replacement battery from the antistatic shipping bag.
   d. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
5. Close the air duct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. If you have not already done so, reinstall the cable management device.
   d. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      > If your system stops at the boot menu, select the option to boot to LOADER.

6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF C800 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations (including ASA configurations), you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick steps - AFF C800

This section gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Access the AFF C800 Installation and Setup Instructions

Video steps - AFF C800

The following video shows how to install and cable your new system.

Animation - Installation and Setup of an AFF C800

If you have a MetroCluster configuration, use the MetroCluster installation content.

MetroCluster Documentation

Detailed steps - AFF C800

This section gives detailed step-by-step instructions for installing an AFF C800 system.

If you have a MetroCluster configuration, use the MetroCluster installation content.

MetroCluster Documentation

Step 1: Prepare for installation

To install your AFF C800 system, you need to create an account and register the system. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the NetApp Hardware Universe (HWU) for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

What you need

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser
    1. Unpack the contents of all boxes.
    2. Record the system serial number from the controllers.

Steps

1. Set up your account:
    a. Log in to your existing account or create an account.
    b. Register (NetApp Product Registration) your system.
2. Download and install NetApp Downloads: Config Advisor on your laptop.
3. Inventory and make a note of the number and types of cables you received.

    The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Connector type	Part number and length	Type of cable…	For…
100 GbE cable	X66211A-05 (112-00595), 0.5m X66211-1 (112-00573), 1m X66211-2 (112-00574), 2m X66211-5 (112-00576), 5m		• HA interconnect • Cluster interconnect network • Storage, Data
10 GbE cable	X6566B-3-R6 (112-00300), 3m; X6566B-5-R6 (112-00301), 5m		• Data
25 GbE cable	X66240A-2 (112-00598), 2m; X66240A-5 (112-00600), 5m		• Data
RJ-45 (order dependent)	Not applicable		• Management
Fibre Channel	X66250-2 (112-00342) 2m; X66250-5 (112-00344) 5m; X66250-15 (112-00346) 15m; X66250-30 (112-00347) 30m		• Network
Micro-USB console cable	Not applicable		• Console connection during software setup
Power cables	Not applicable		Connecting the PSUs to power source

4. Download and complete the Cluster Configuration Worksheet.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.

   SuperRail kit installation instructions

2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

3. Attach cable management devices (as shown).

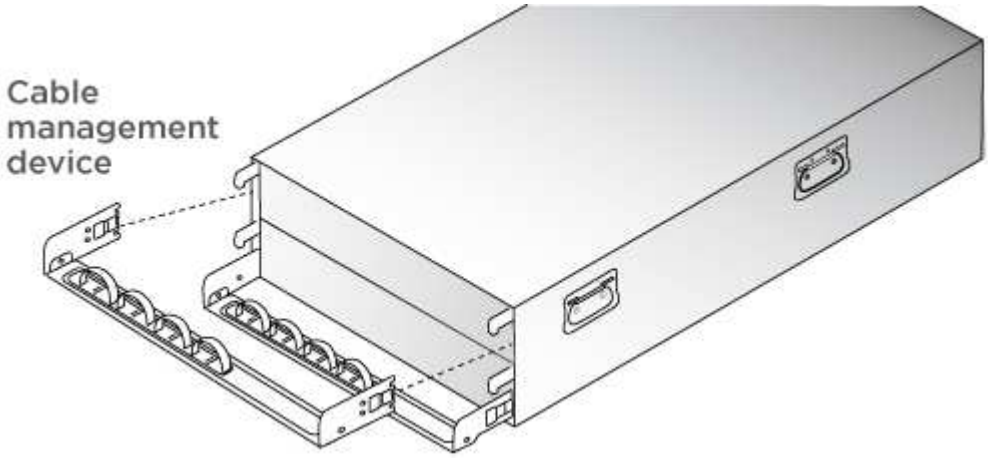

4. Place the bezel on the front of the system.

Step 3: Cable controllers

There is required cabling for your platform's cluster using the two-node switchless cluster method or the cluster interconnect network method. There is optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cable to a host network and storage.

Required cabling: Cable controllers to a cluster

Cable the controllers to a cluster by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Cable a two-node switchless cluster

Management network ports on the controllers are connected to switches. The HA interconnect and cluster interconnect ports are cabled on both controllers.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

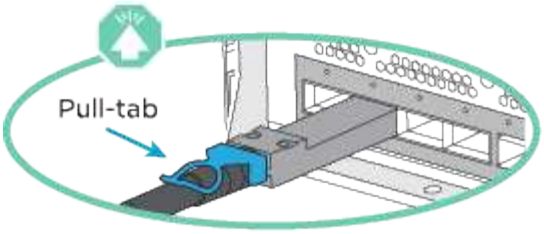

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or the tabulated steps to complete the cabling between the controllers and the switches:

Animation - Cable a two-node switchless cluster

Step	Perform on each controller module
1	Cable the HA interconnect ports: • e0b to e0b • e1b to e1b 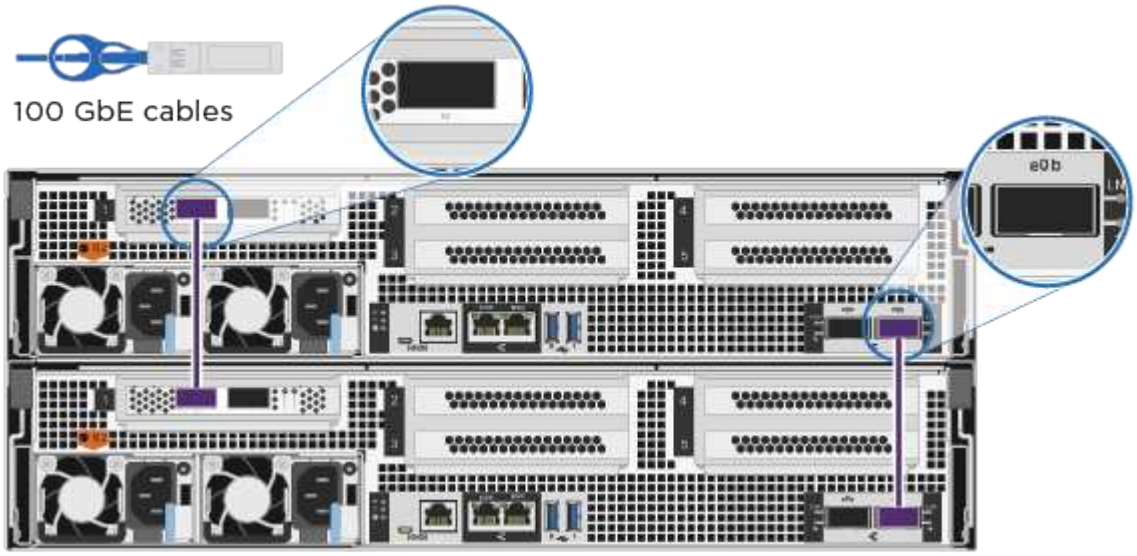

Step	Perform on each controller module
2	Cable the cluster interconnect ports: • e0a to e0a • e1a to e1a 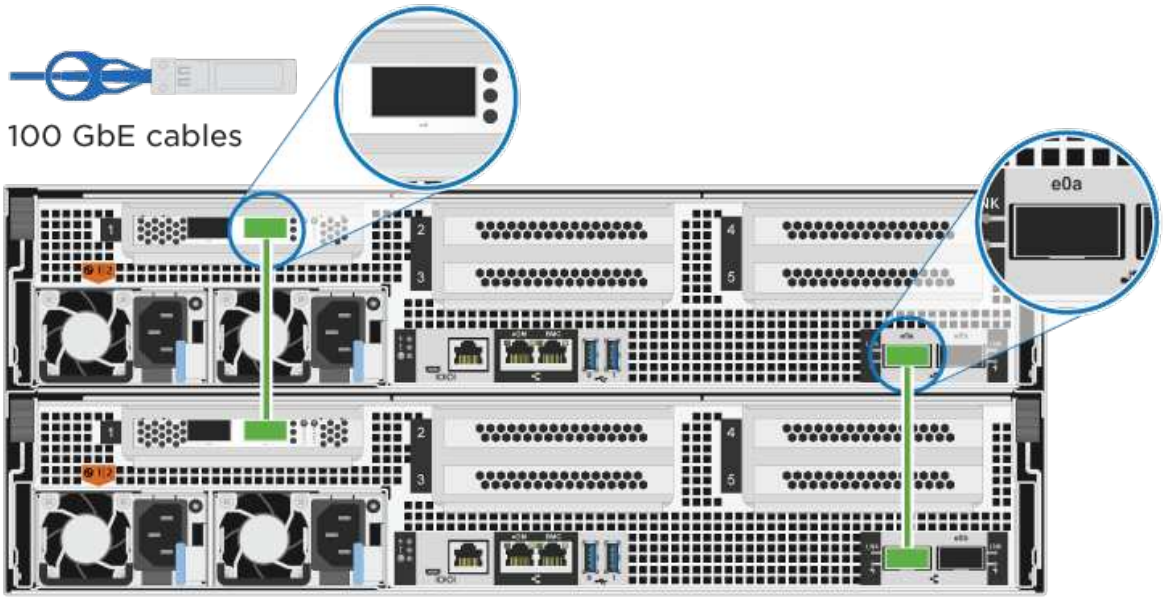
3	Cable the management ports to the management network switches 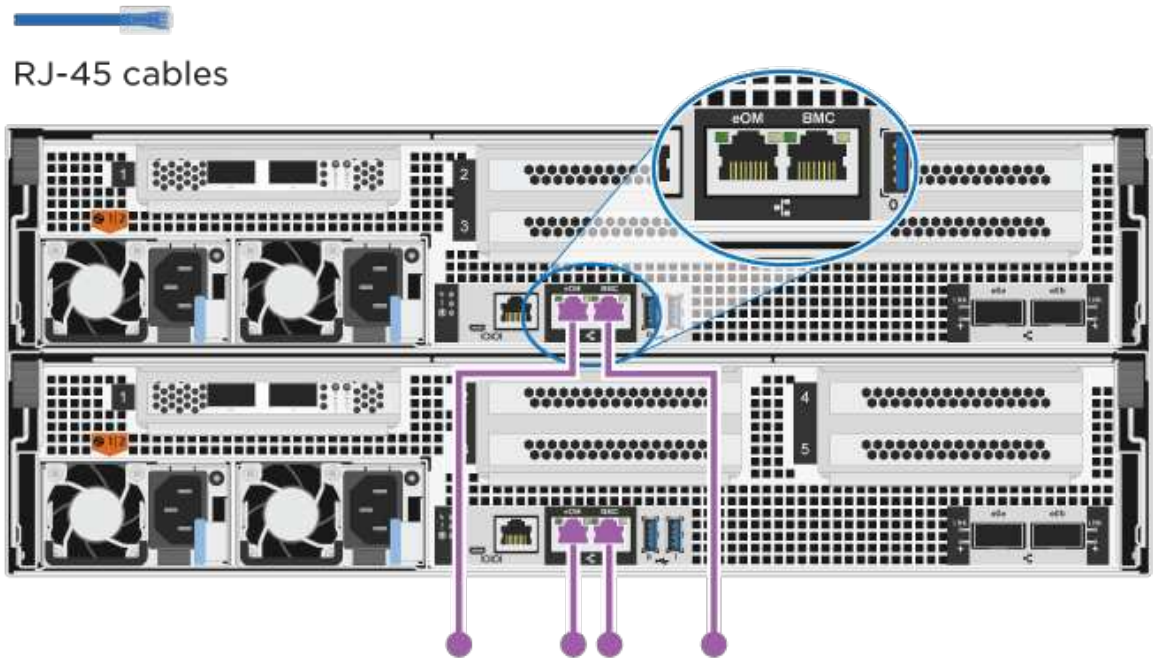
!	DO NOT plug in the power cords at this point.

2. To perform optional cabling, see:
   - Option 1: Cable to a Fibre Channel host network
   - Option 2: Cable to a 10GbE host network
   - Option 3: Cable the controllers to a single drive shelf
   - Option 4: Cable the controllers to two drive shelves
3. To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 2: Cable a switched cluster

Cluster interconnect and management network ports on the controllers are connected to switches while the HA interconnect ports are cabled on both controllers.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

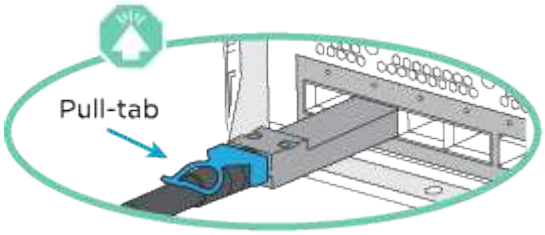

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or the tabulated steps to complete the cabling between the controllers and the switches:

   Animation - Cable a switched cluster

Step	Perform on each controller module
1	Cable the HA interconnect ports: • e0b to e0b • e1b to e1b 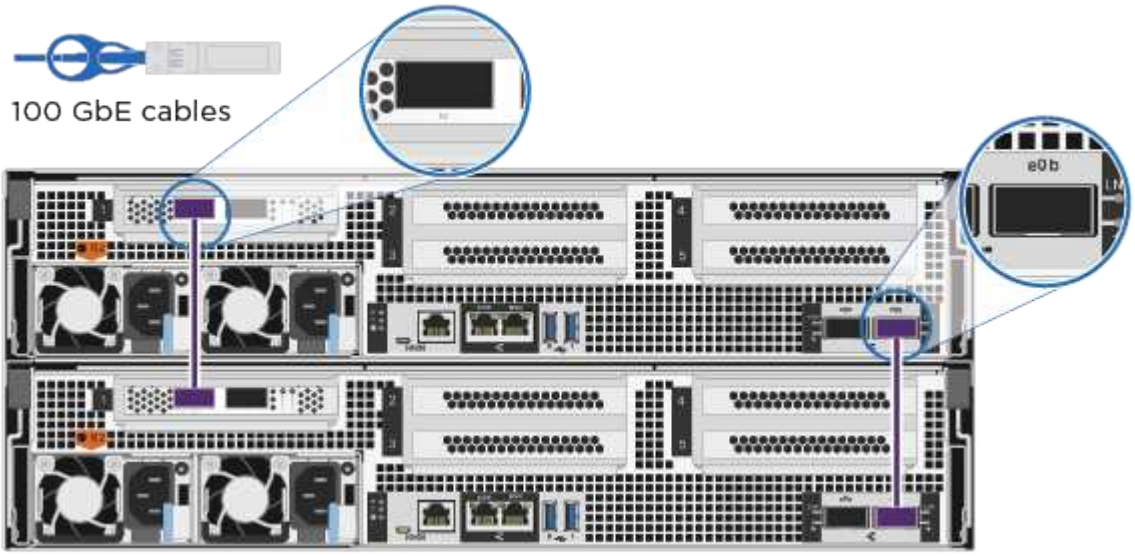
2	Cable the cluster interconnect ports to the 100 GbE cluster interconnect switches. • e0a • e1a 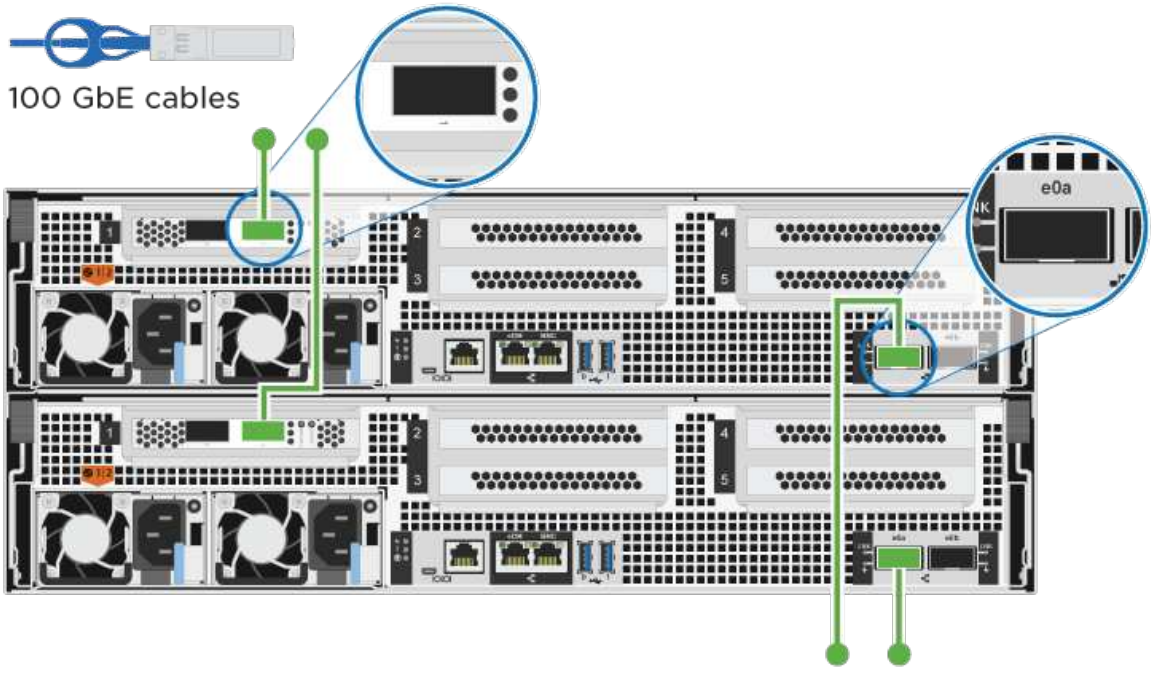

Step	Perform on each controller module
3	Cable the management ports to the management network switches 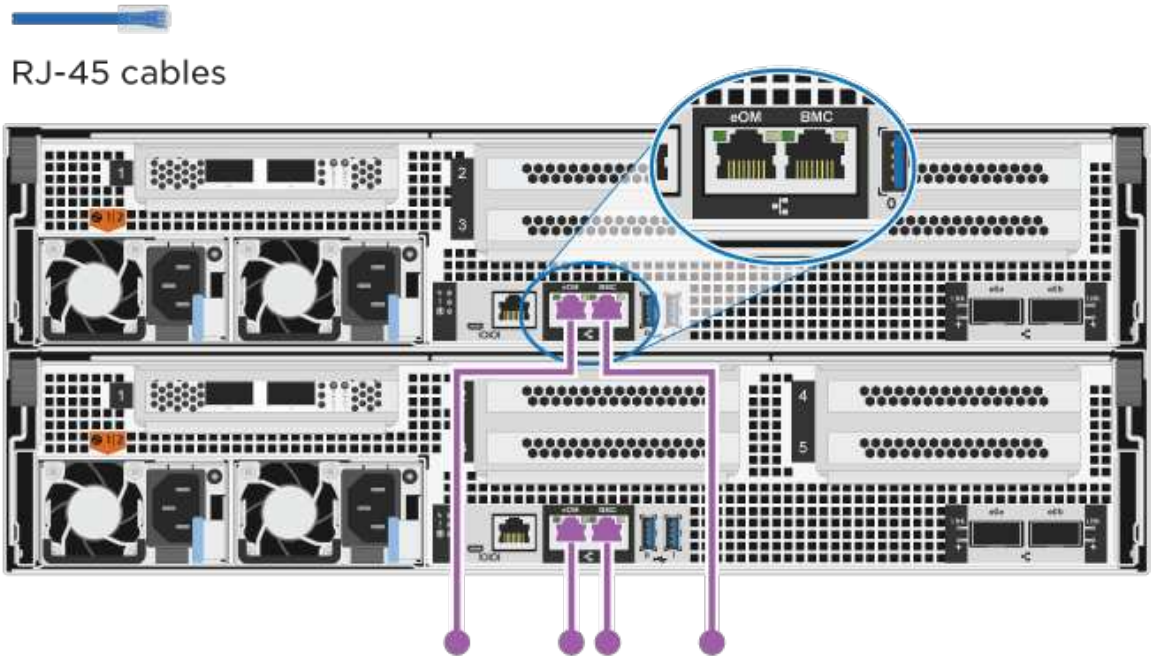
	DO NOT plug in the power cords at this point.

2. To perform optional cabling, see:
   - Option 1: Cable to a Fibre Channel host network
   - Option 2: Cable to a 10GbE host network
   - Option 3: Cable the controllers to a single drive shelf
   - Option 4: Cable the controllers to two drive shelves
3. To complete setting up your system, see Step 4: Complete system setup and configuration.

Optional cabling: Cable configuration-dependent options

You have configuration-dependent optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cabling to a host network and storage.

Option 1: Cable to a Fibre Channel host network

Fibre Channel ports on the controllers are connected to Fibre Channel host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

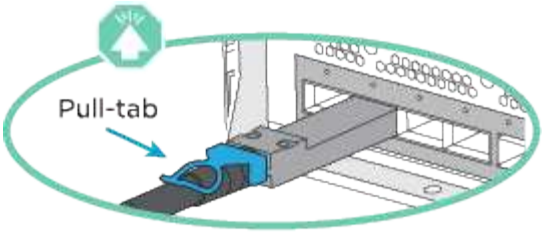

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Step	Perform on each controller module
1	Cable ports 2a through 2d to the FC host switches. 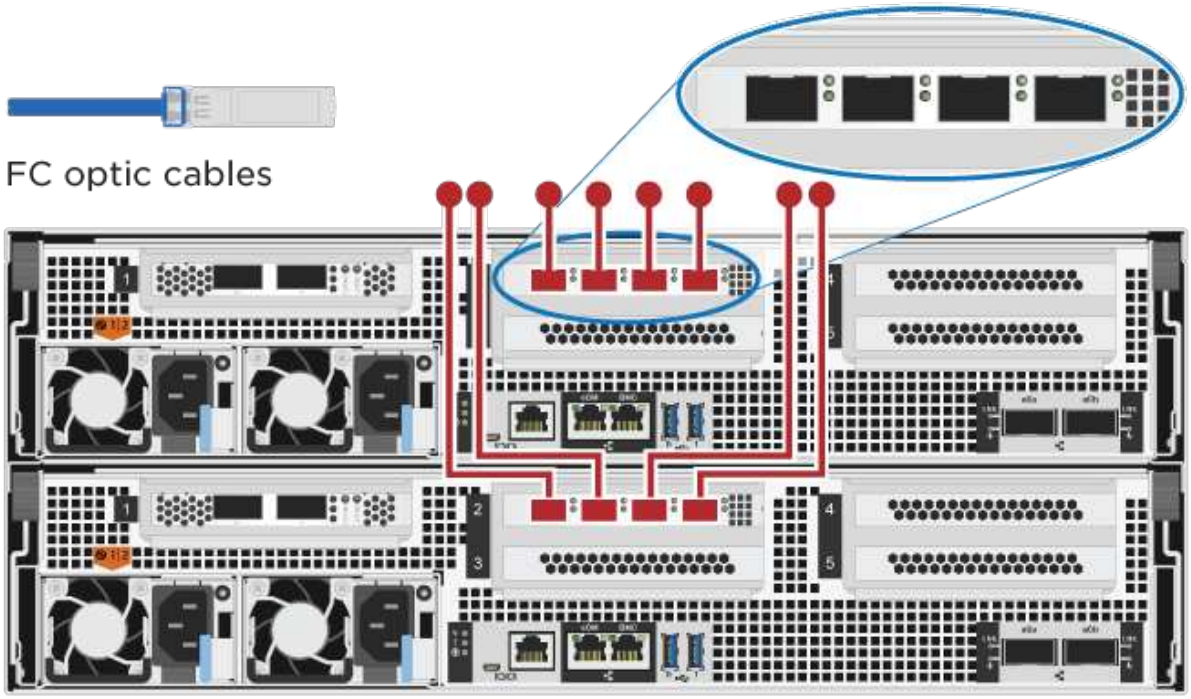
2	To perform other optional cabling, choose from: • Option 3: Cable the controllers to a single drive shelf • Option 4: Cable the controllers to two drive shelves
3	To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 2: Cable to a 10GbE host network

10GbE ports on the controllers are connected to 10GbE host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

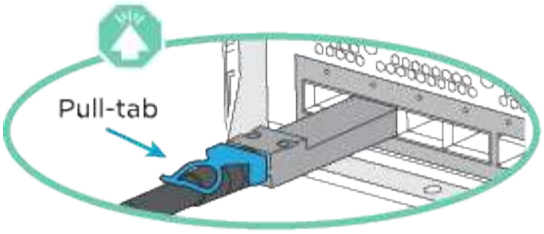

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Step	Perform on each controller module
1	Cable ports e4a through e4d to the 10GbE host network switches. 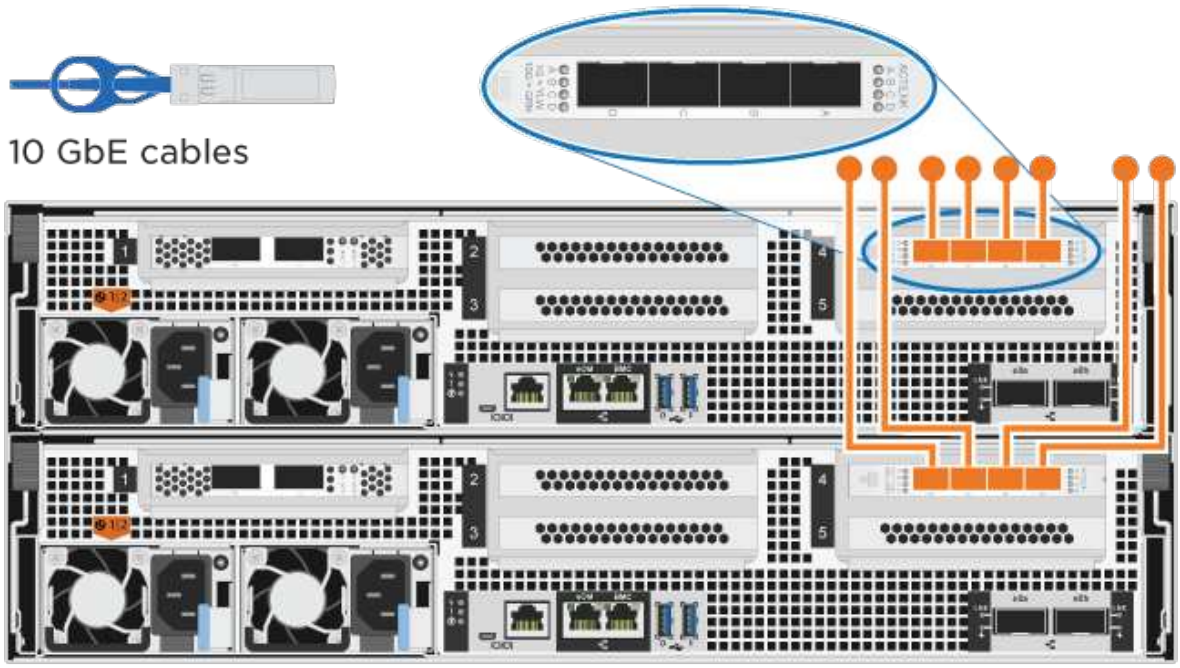
2	To perform other optional cabling, choose from: • Option 3: Cable the controllers to a single drive shelf • Option 4: Cable the controllers to two drive shelves
3	To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 3: Cable the controllers to a single drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

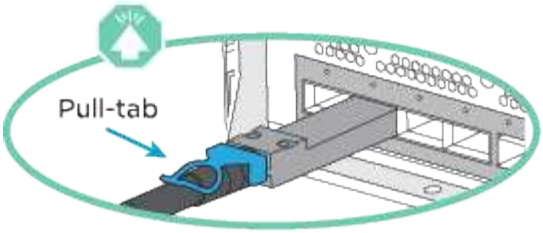

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to cable your controllers to a single shelf:

Animation - Cable the controllers to a single drive shelf

Step	Perform on each controller module
1	Cable controller A to the shelf: 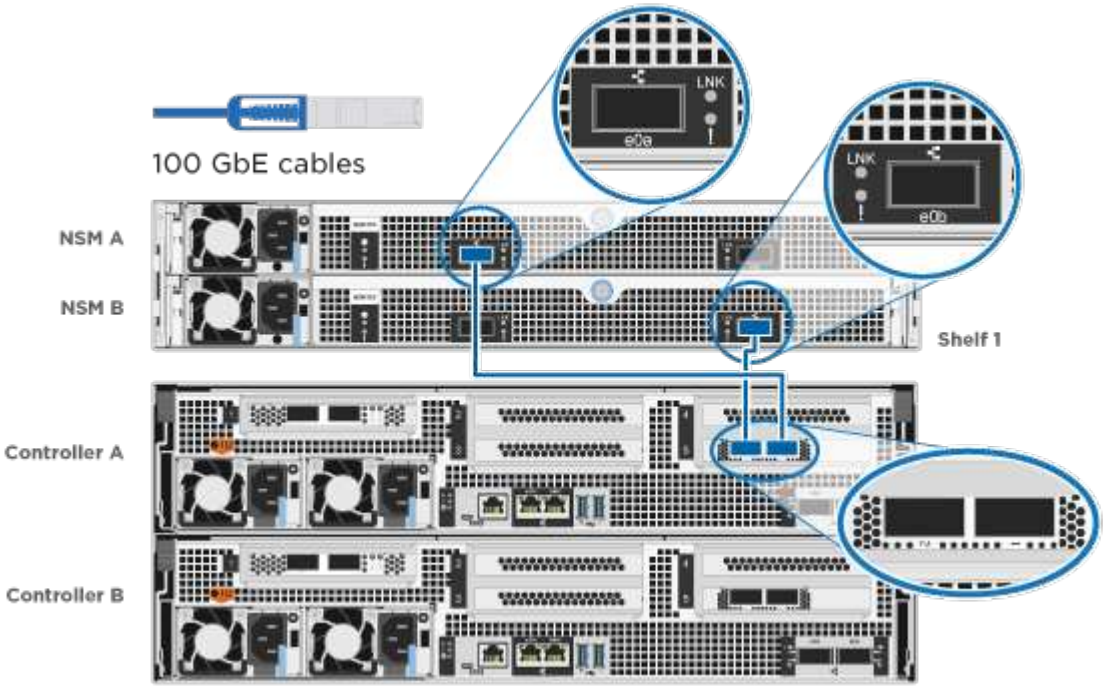

Step	Perform on each controller module
2	Cable controller B to the shelf: 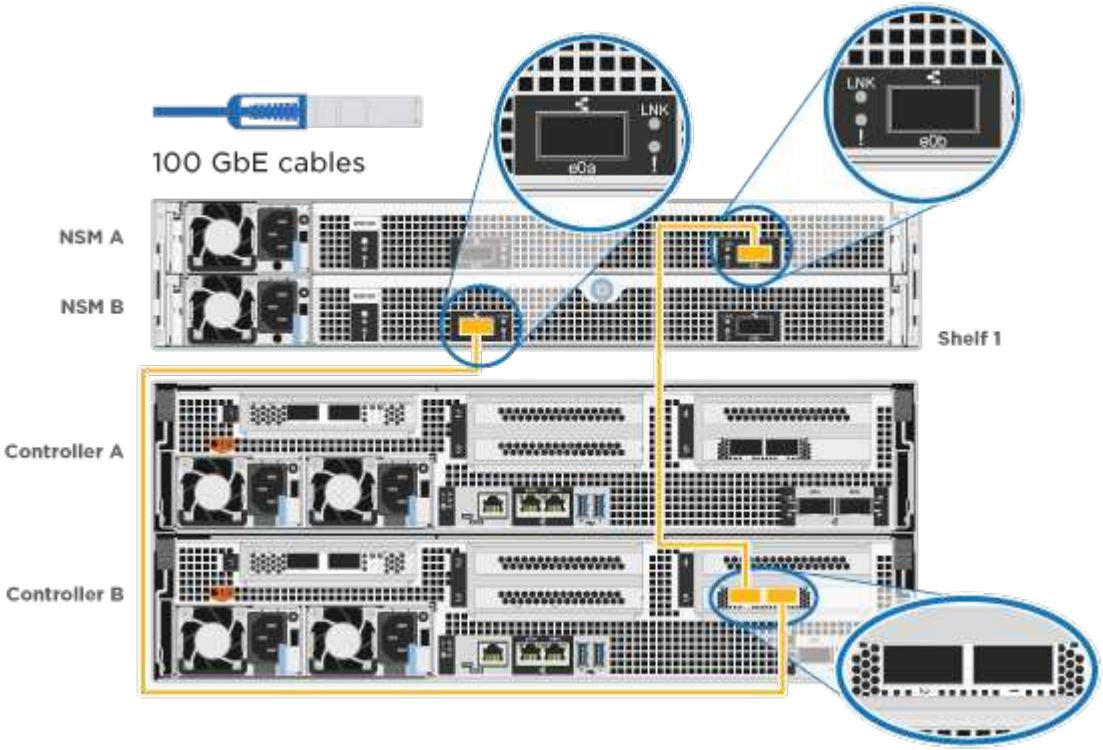

To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 4: Cable the controllers to two drive shelves

You must cable each controller to the NSM modules on both NS224 drive shelves.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

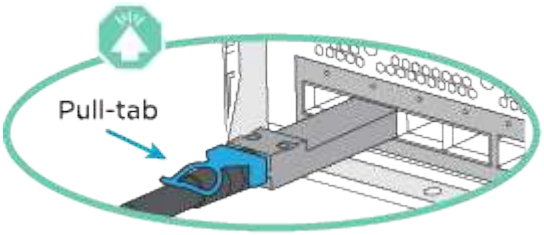

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to cable your controllers to two drive shelves:

Animation - Cable the controllers to two drive shelves

Step	Perform on each controller module
1	Cable controller A to the shelves: 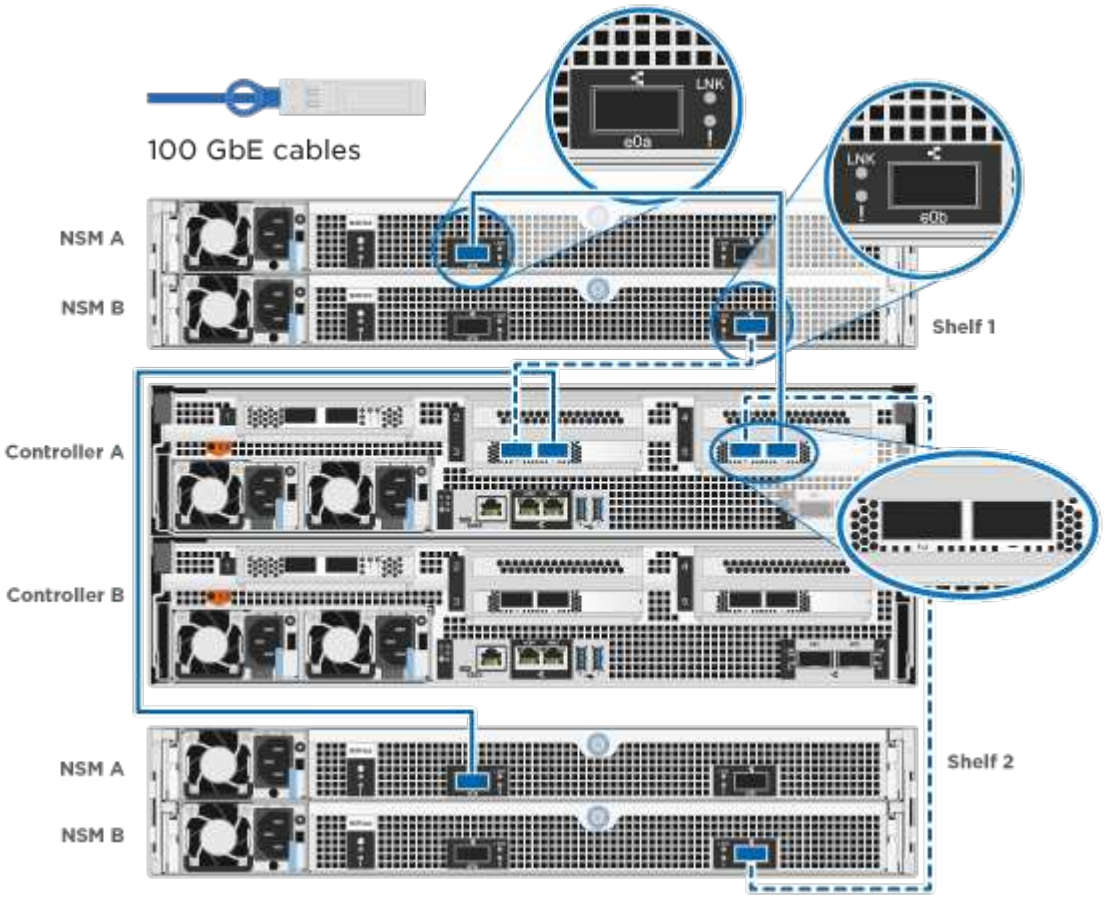

Step	Perform on each controller module
2	Cable controller B to the shelves: 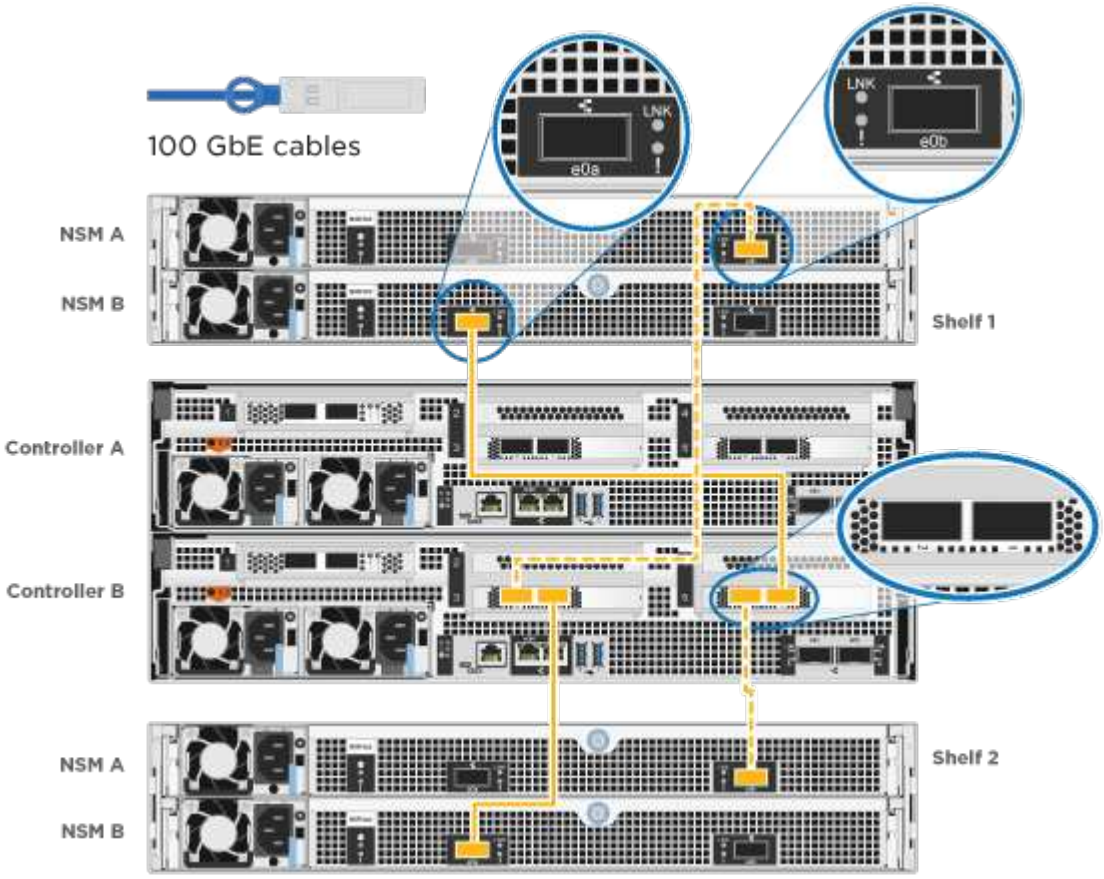

To complete setting up your system, see Step 4: Complete system setup and configuration.

Step 4: Complete system setup and configuration

Complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Complete system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Use the following animation to power on and set shelf IDs for one or more drive shelves:

   For NS224 drive shelves, shelf IDs are pre-set to 00 and 01. If you want to change the shelf IDs, use the straightened end of a paperclip, or narrow tipped ball point pen to access the shelf ID button behind the faceplate.

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

The system begins to boot. Initial booting may take up to eight minutes.

3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Connect your laptop to the Management switch:

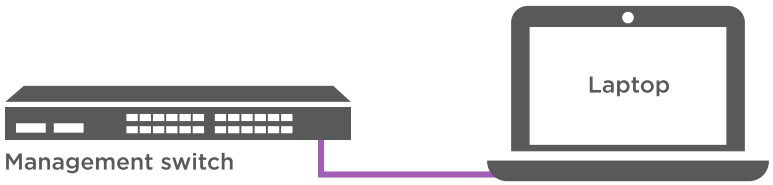

5. Select an ONTAP icon listed to discover:

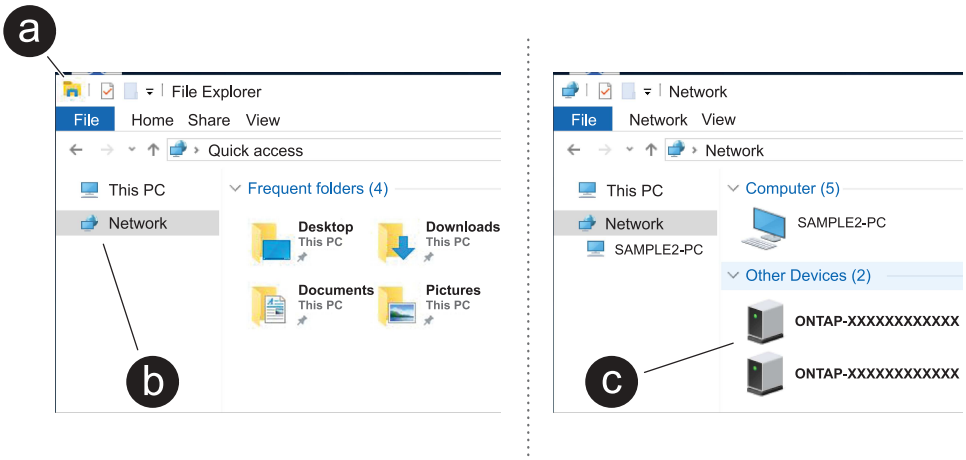

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

   System Manager opens.

6. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Complete system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your system.

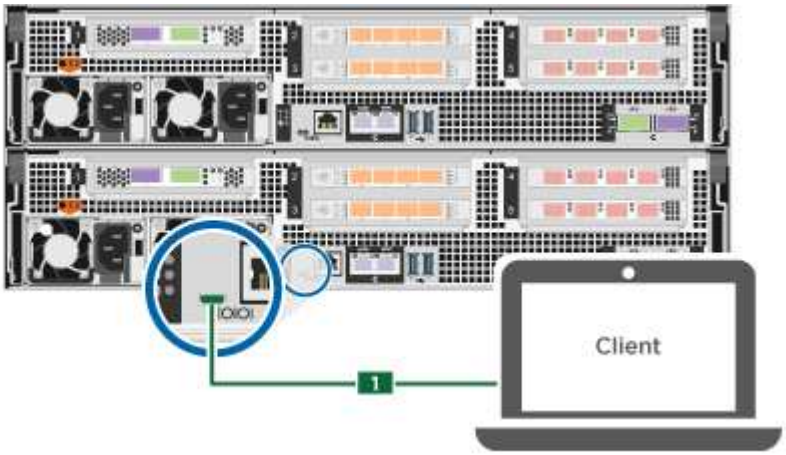

   c. Connect the laptop or console to the switch on the management subnet.

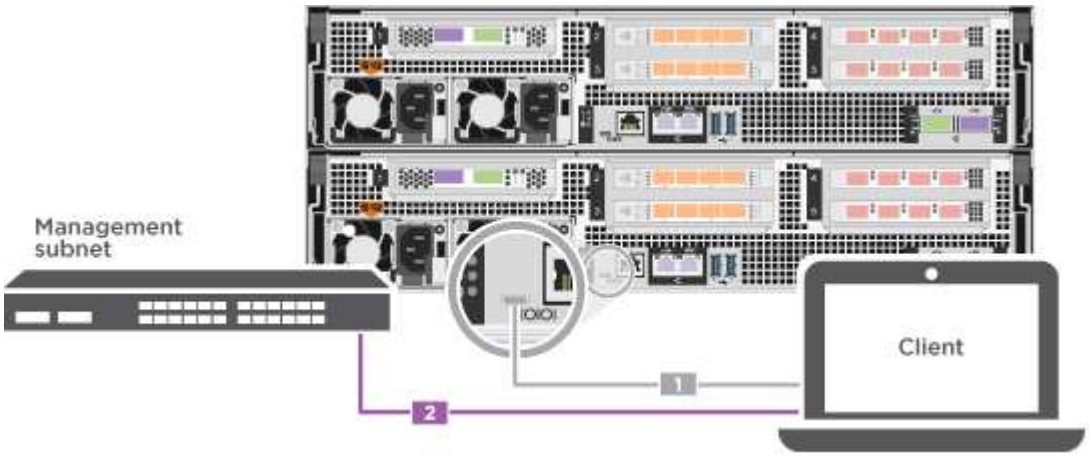

   d. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Use the following animation to power on and set shelf IDs for one or more drive shelves:

For NS224 drive shelves, shelf IDs are pre-set to 00 and 01. If you want to change the shelf IDs, use the straightened end of a paperclip, or narrow tipped ball point pen to access the shelf ID button behind the faceplate.

Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

4. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the ONTAP Configuration Guide.
6. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain AFF C800 hardware

Maintain the hardware of your AFF C800 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the AFF C800 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF C800 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Drive	A drive is a device that provides the physical storage media for data.
Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.
NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.

PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - AFF C800

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your AFF C800 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - AFF C800

Before replacing the boot media in your AFF C800, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the

correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - AFF C800

Shut down the impaired controller in your AFF C800 storage system to prevent data loss nd maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a

healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   `cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - AFF C800

The boot media in your AFF C800 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

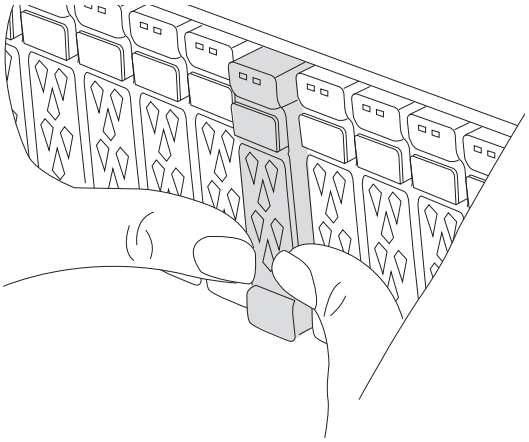

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

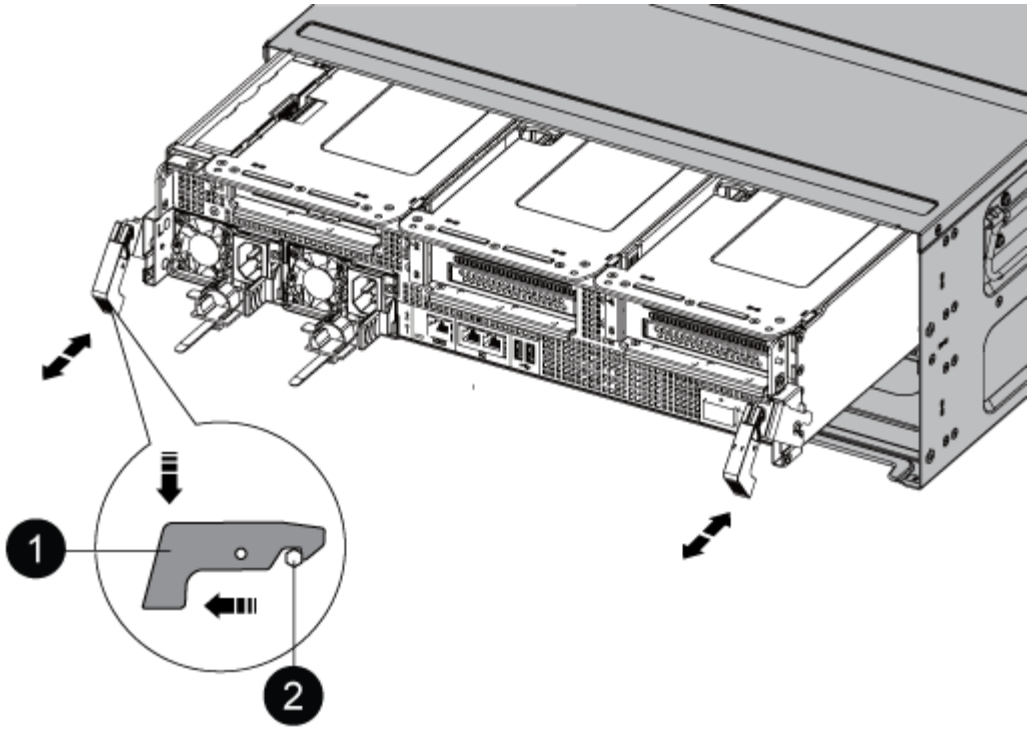

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

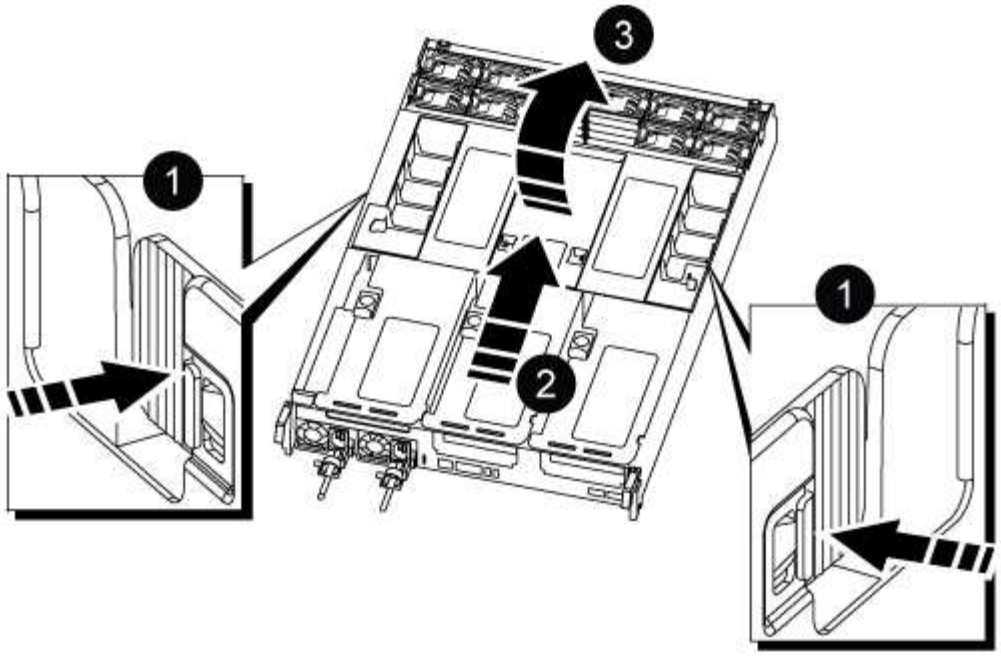

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

10. Locate the boot media in the controller module and replace it:

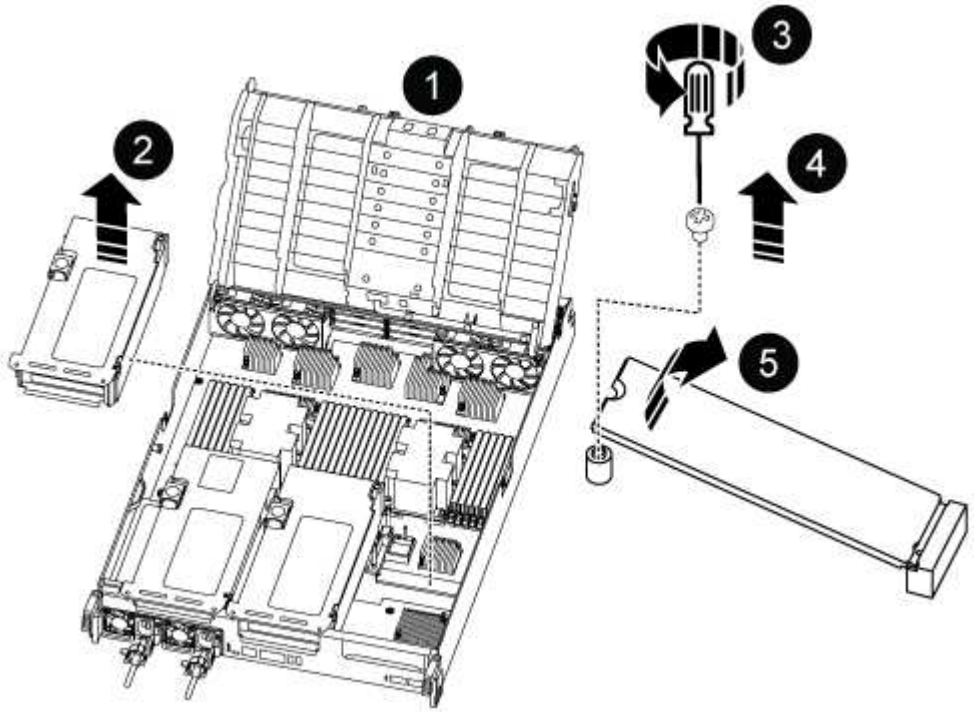

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.

b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.

11. Install the replacement boot media into the controller module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down toward the motherboard.

c. Secure the boot media to the motherboard using the boot media screw.

Do not over-tighten the screw or you might damage the boot media.

12. Reinstall the riser into the controller module.

13. Close the air duct:

a. Rotate the air duct downward.

b. Slide the air duct toward the risers until it clicks into place.

14. Install the controller module:

a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module half-way into the way into the system.

b. Recable the controller module, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

The controller module begins to boot and stops at the LOADER prompt.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - AFF C800

After installing the new boot media device in your AFF C800 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and

determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

   Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

   Show example of configuration error finding prompts

   ```
   Error when fetching key manager config from partner ${partner_ip}:
   ${status}

   Has key manager been configured on this system

   Is the key manager onboard
   ```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - AFF C800

If a component in your AFF C800 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - AFF C800

Get started with replacing the boot media in your AFF C800 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - AFF C800

Before replacing the boot media in your AFF C800 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - AFF C800

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - AFF C800

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: System is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media and prepare for manual boot recovery - AFF C800

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

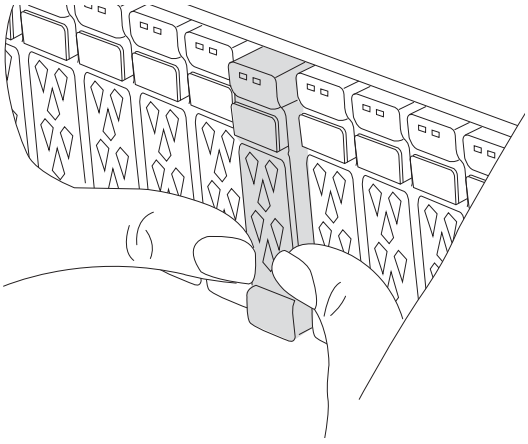

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management

device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

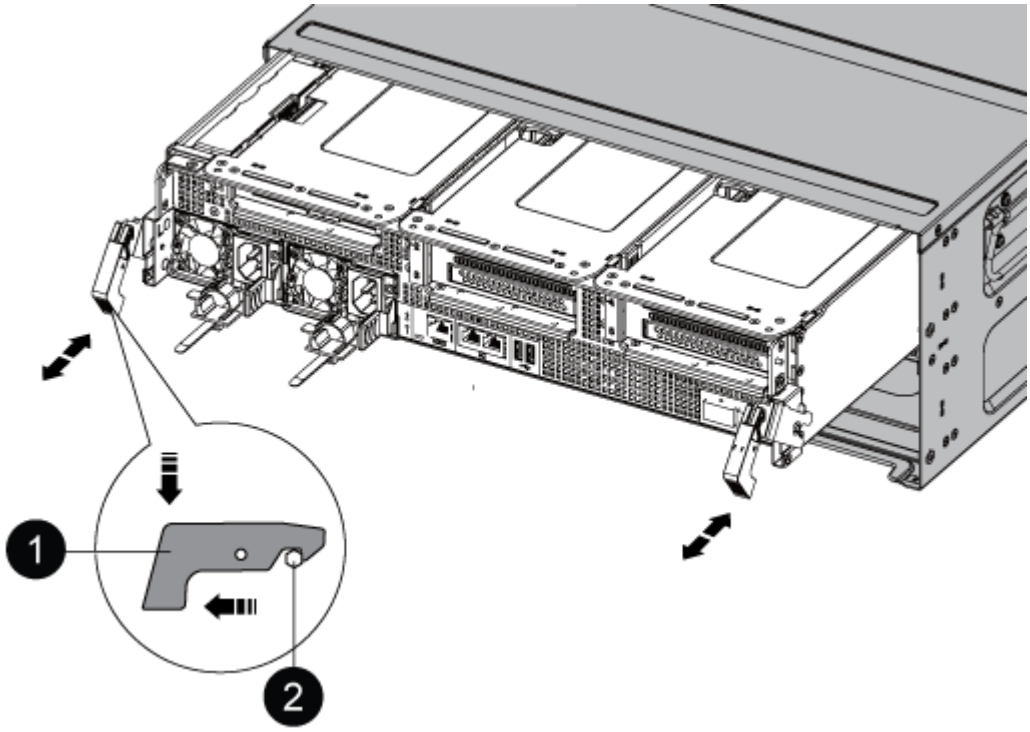

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

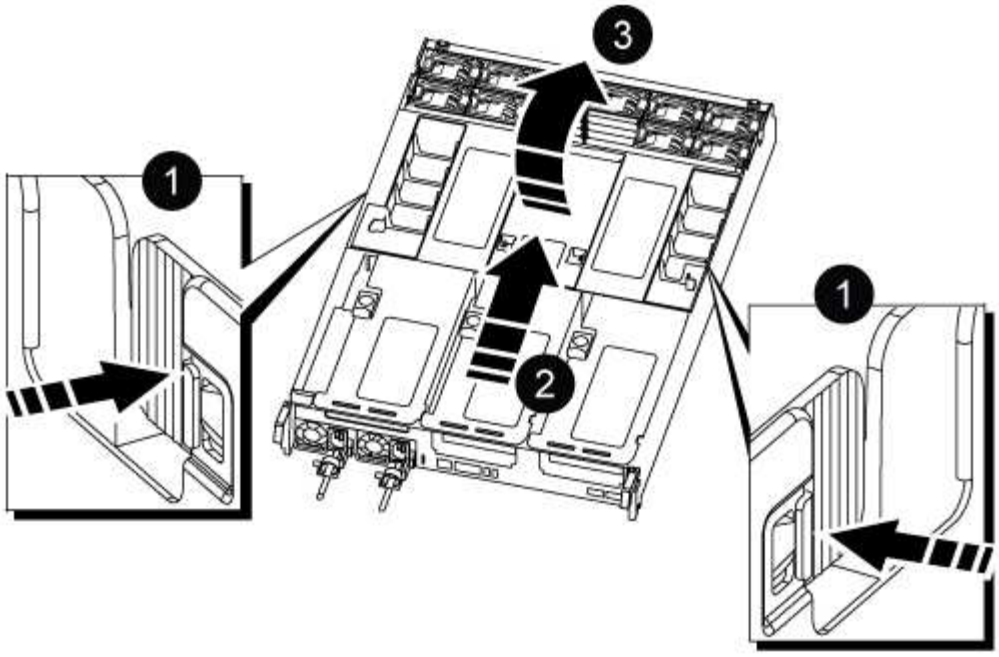

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 2: Replace the boot media

You locate the failed boot media in the controller module by removing Riser 3 on the controller module before you can replace the boot media.

You need a Phillips head screwdriver to remove the screw that holds the boot media in place.

1. Locate the boot media:

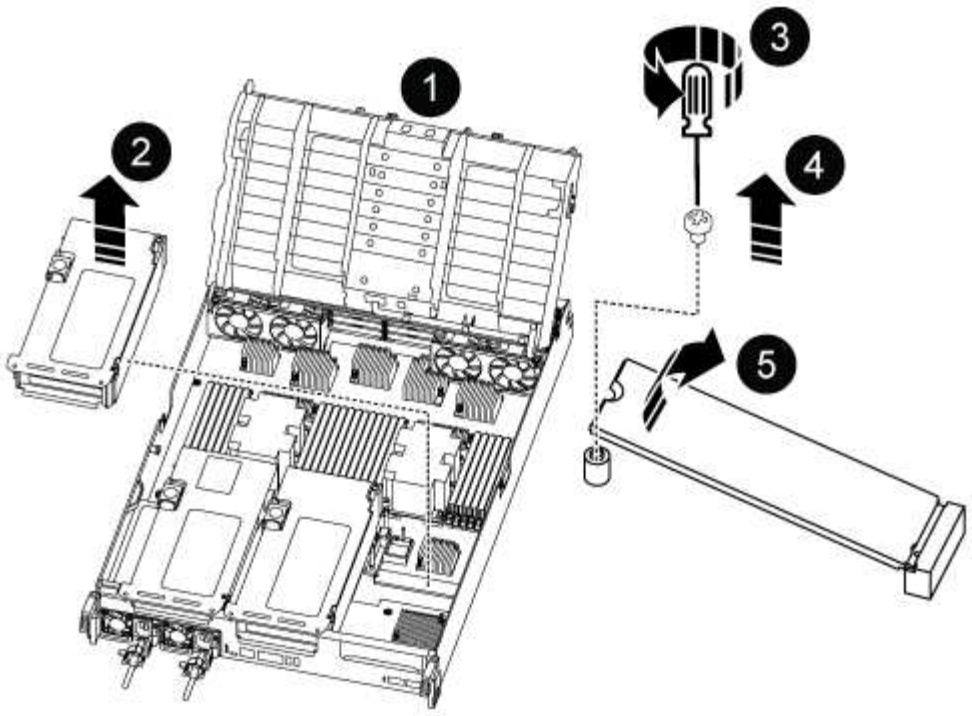

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

2. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.
   b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.
3. Install the replacement boot media into the controller module:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the motherboard.
   c. Secure the boot media to the motherboard using the boot media screw.

      Do not over-tighten the screw or you might damage the boot media.
4. Reinstall the riser into the controller module.

5. Close the air duct:
    a. Rotate the air duct downward.
    b. Slide the air duct toward the risers until it clicks into place.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed is without a boot image so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
    - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
    - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
    a. Download the service image to your work space on your laptop.
    b. Unzip the service image.

        If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

        There are two folders in the unzipped service image file:

        - boot
        - efi

    c. Copy the efi folder to the top directory on the USB flash drive.

        If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

    The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

    a. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct:
    a. Swing the air duct all the way down to the controller module.
    b. Slide the air duct toward the risers until the locking tabs click into place.

c. Inspect the air duct to make sure that it is properly seated and locked into place.

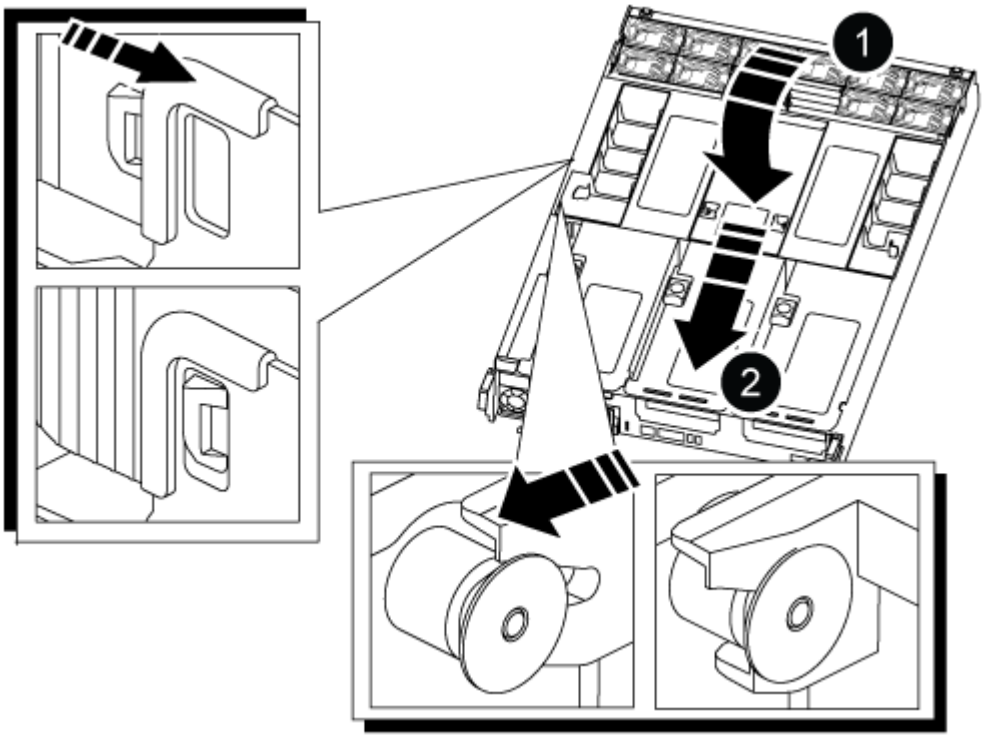

1	Air duct
2	Risers

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.

5. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

6. Gently push the controller module all the way into the system until the controller module locking hooks begin to rise, firmly push on the locking hooks to finish seating the controller module, and then swing the locking hooks into the locked position over the pins on the controller module.
7. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the

controller to boot to LOADER.

Manual boot media recovery from a USB drive - AFF C800

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore encryption - AFF C800

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - AFF C800

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Replace the chassis - AFF C800

To replace the chassis, you must move the bezel, controller modules, and NVMe drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in its place.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF C800

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - AFF C800

Move the power supplies, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

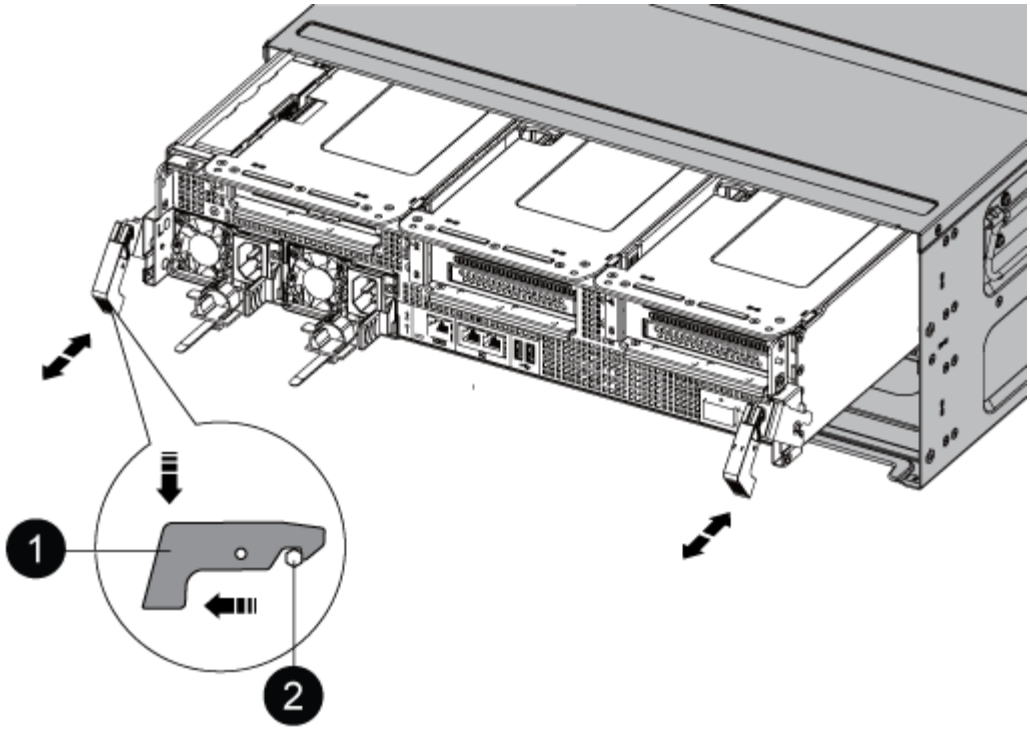

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.

3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.

5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process by pressing `Ctrl-C`.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - AFF C800

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - AFF C800

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Do not downgrade the BIOS version of the *replacement* controller to match the partner controller or the old controller module.

Shut down the impaired controller - AFF C800

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   `cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - AFF C800

To replace the controller, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

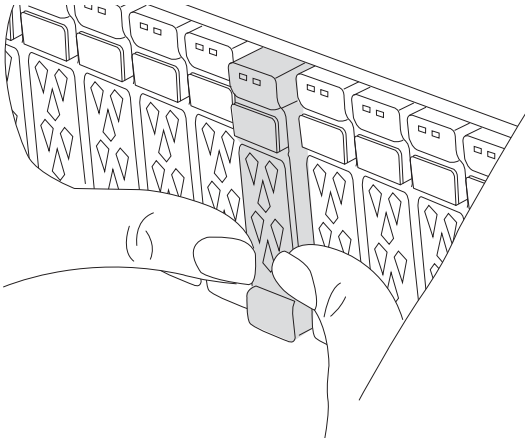

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

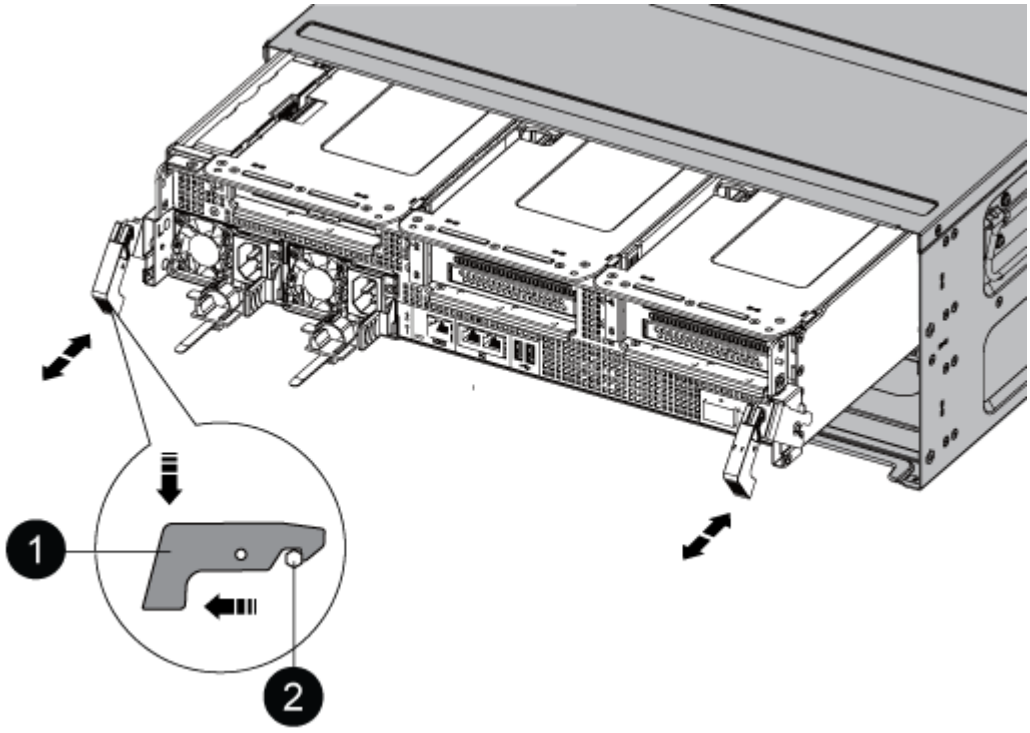

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
   .

9. Place the controller module on a stable, flat surface.
10. Open the controller module air duct:
    a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
    b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

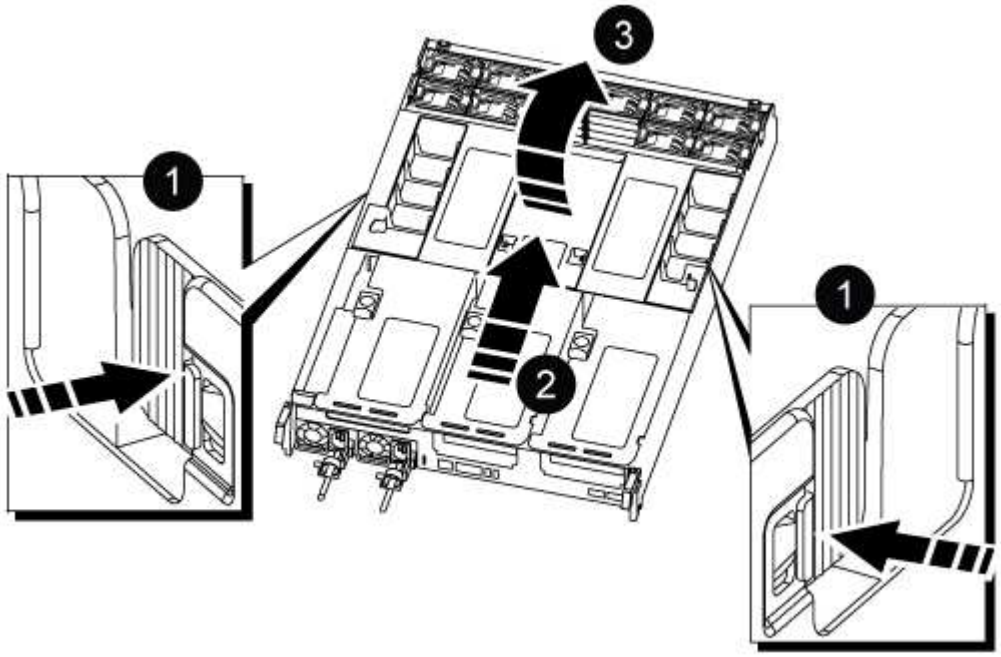

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 2: Move the power supplies

You must move the power supplies from the impaired controller module to the replacement controller module when you replace a controller module.

1. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

> The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

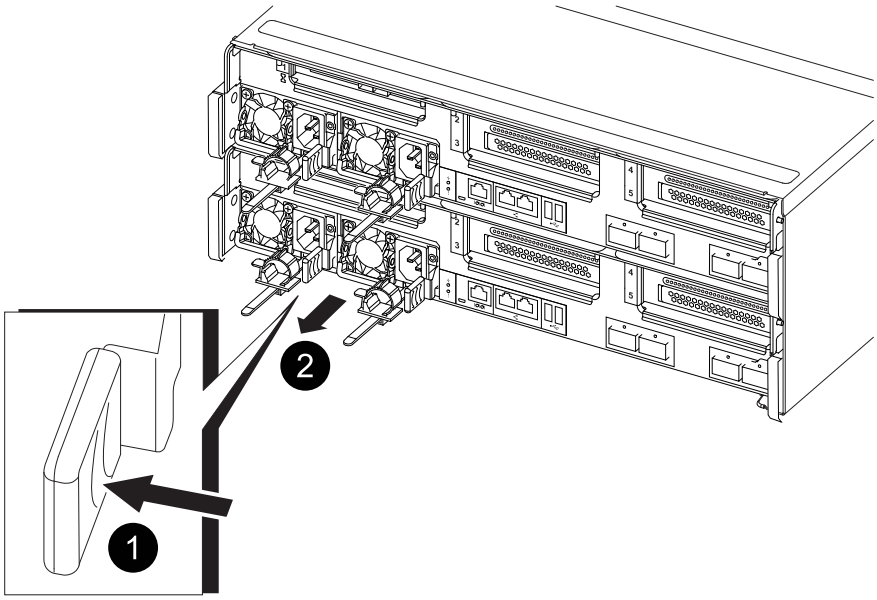

1	Blue power supply locking tab
2	Power supply

2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

> To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

You must move the fans from the impaired controller module to the replacement module when replacing a failed controller module.

1. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

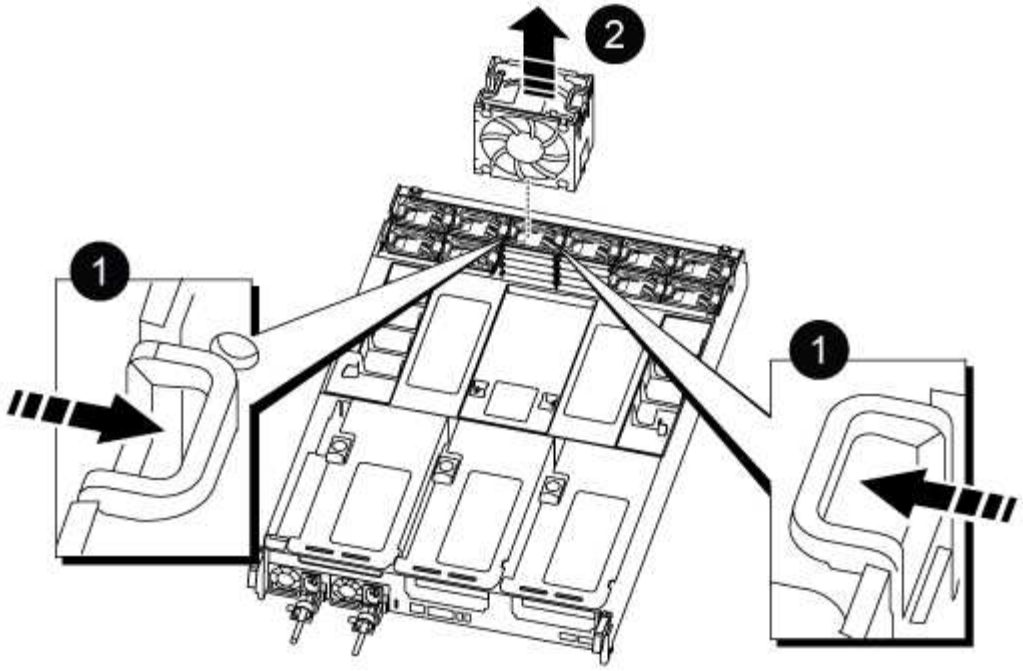

1	Fan locking tabs
2	Fan module

2. Move the fan module to the replacement controller module, and then install the fan module by aligning its edges with the opening in the controller module, and then sliding the fan module into the controller module until the locking latches click into place.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the NVDIMM battery

When replacing the controller module, you must move the NVRAM battery from the impaired controller module to the replacement controller module

1. Open the air duct cover and locate the NVDIMM battery in the riser.

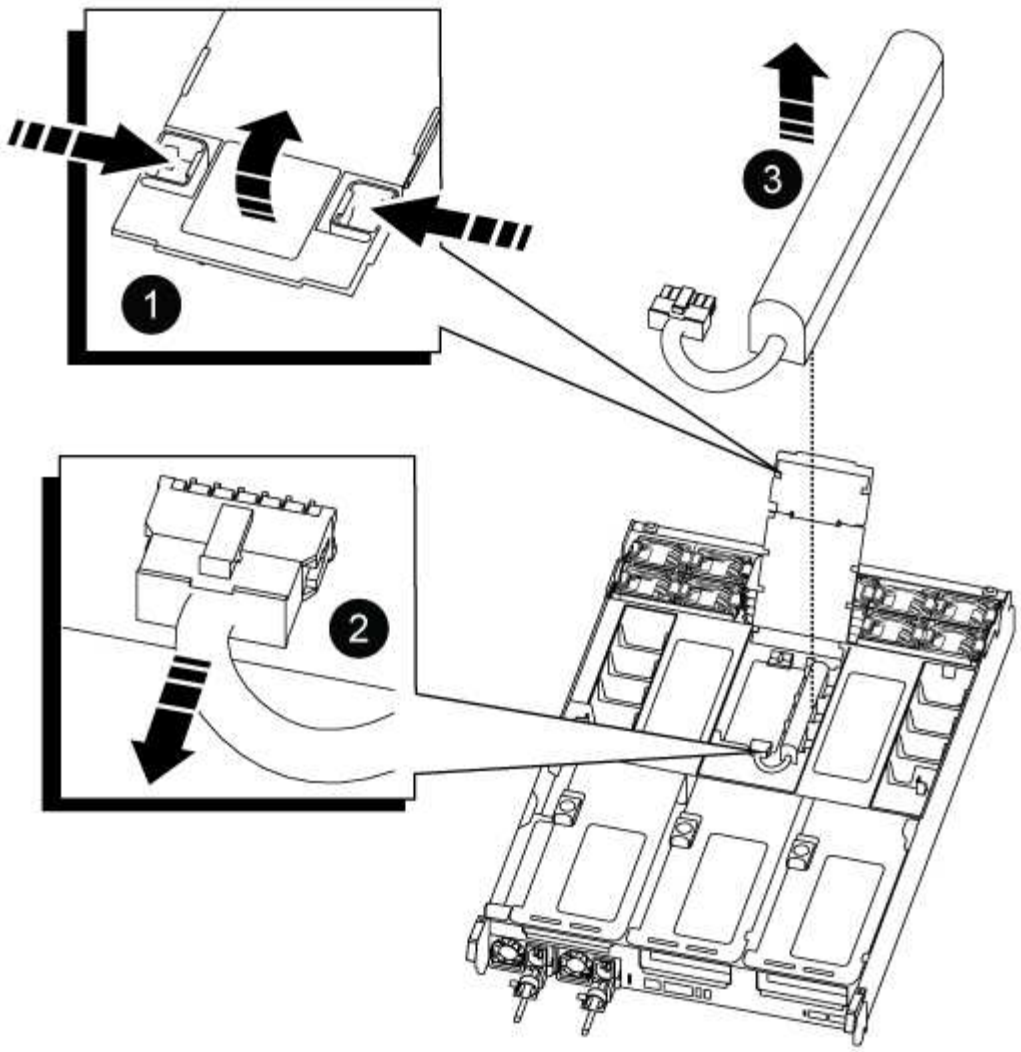

1	Air duct riser
2	NVDIMM battery plug
3	NVDIMM battery pack

Attention: The NVDIMM battery control board LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
3. Grasp the battery and lift the battery out of the air duct and controller module.
4. Move the battery pack to the replacement controller module and then install it in the NVDIMM air duct:
   a. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   b. Plug the battery plug into the riser socket and make sure that the plug locks into place.

Step 5: Remove the PCIe risers

As part of the controller replacement process, you must remove the PCIe modules from the impaired controller module. You must install them into the same location in the replacement controller module once the NVDIMMS and DIMMs have moved to the replacement controller module.

1. Remove the PCIe riser from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

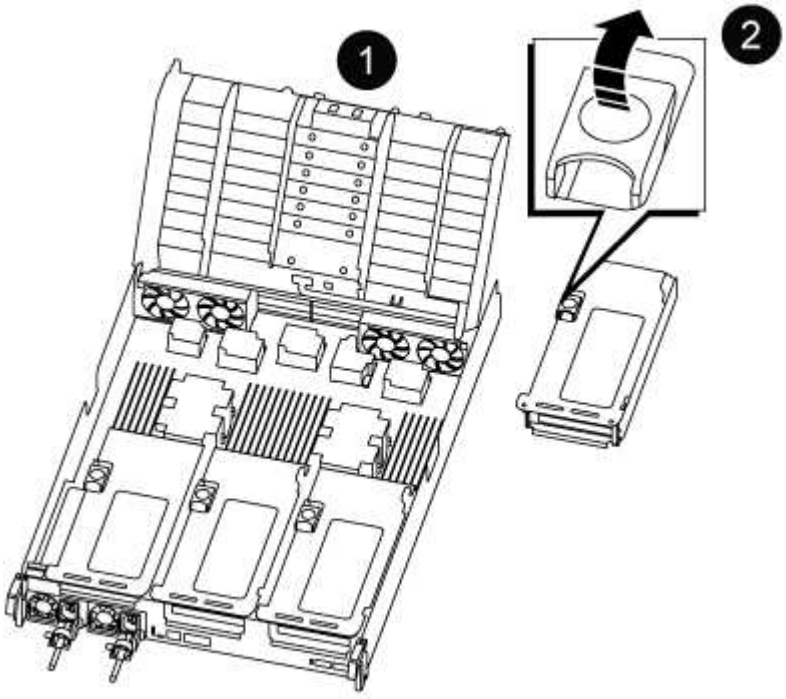

1	Air duct
2	Riser 1 (left riser), Riser 2 (middle riser), and 3 (right riser) locking latches

2. Repeat the preceding step for the remaining risers in the impaired controller module.
3. Repeat the above steps with the empty risers in the replacement controller and put them away.

Step 6: Move system DIMMs

To move the DIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

1. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.

2. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

3. Locate the slot where you are installing the DIMM.
4. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

5. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
6. Repeat these steps for the remaining DIMMs.

Step 7: Move the NVDIMMs

To move the NVDIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

1. Locate the NVDIMMs on your controller module.

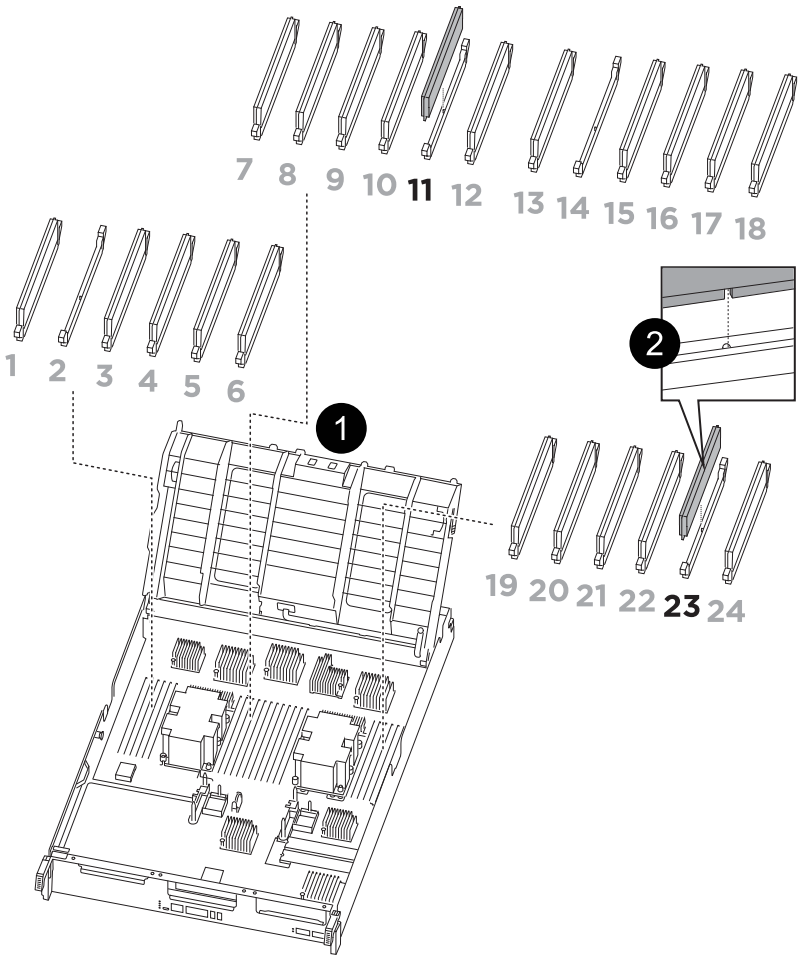

- NVDIMM: SLOTS 11 & 23

1	Air duct
2	NVDIMMs

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

4. Locate the slot where you are installing the NVDIMM.
5. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
7. Repeat the preceding steps to move the other NVDIMM.

Step 8: Move the boot media

You must move the boot media device from the impaired controller and install it in the replacement controller.

The boot media is located under Riser 3.

1. Locate the boot media:

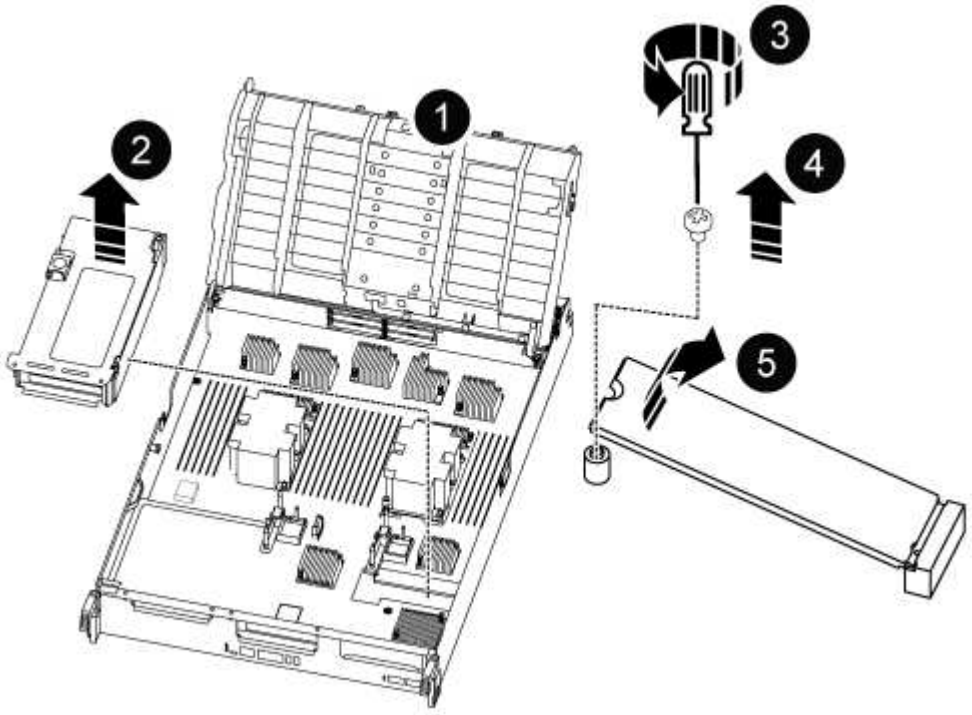

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

2. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.

b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.

3. Move the boot media to the new controller module and install it:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the motherboard.
   c. Secure the boot media to the motherboard using the boot media screw.

      Do not over-tighten the screw or you might damage the boot media.

Step 9: Install the PCIe risers

You install the PCIe risers in the replacement controller module after moving the DIMMs, NVDIMMs, and boot media.

1. Install the riser into the replacement controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

   d. Reinsert any SFP or QSFP modules that were removed from the PCIe cards.
2. Repeat the preceding step for the remaining PCIe risers.

Step 10: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

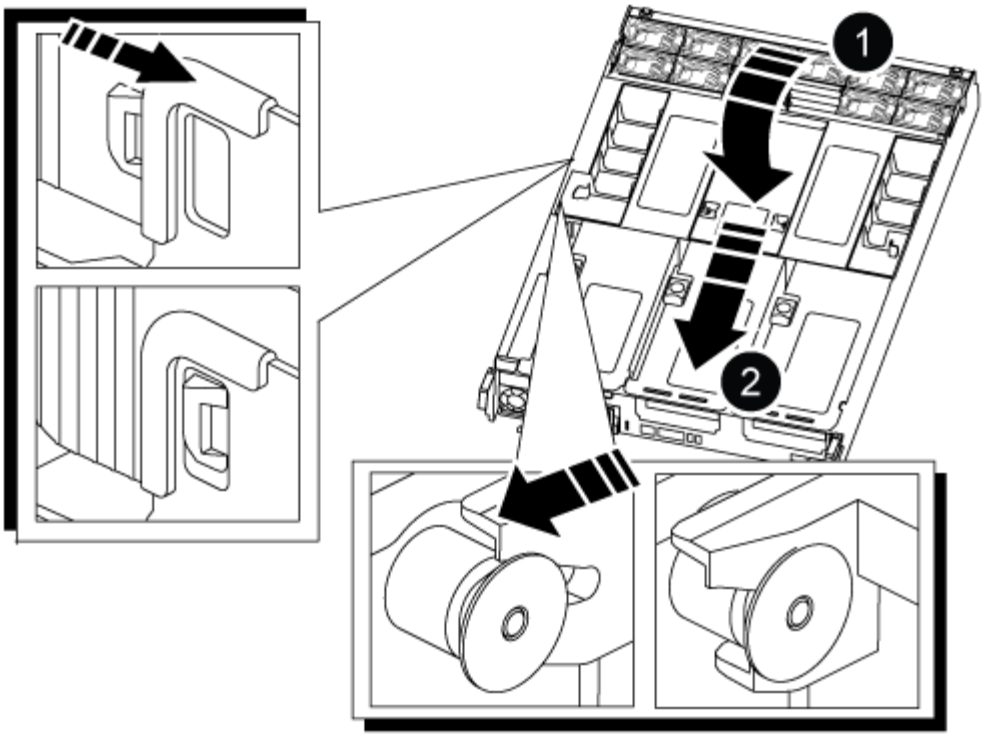

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:

   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

5. Plug the system cables and transceiver modules into the controller module and reinstall the cable management device.

6. Plug the power cables into the power supplies and reinstall the power cable retainers.

   The controller module begins to boot as soon as it is connected to power. Be prepared to interrupt the boot process.

   If your system has DC power supplies, make sure the thumbscrews on the power supply cable are tight.

Restore and verify the system configuration - AFF C800

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the chassis

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA`

state: `ha-config show`

The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF C800

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and

then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF C800

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF C800

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

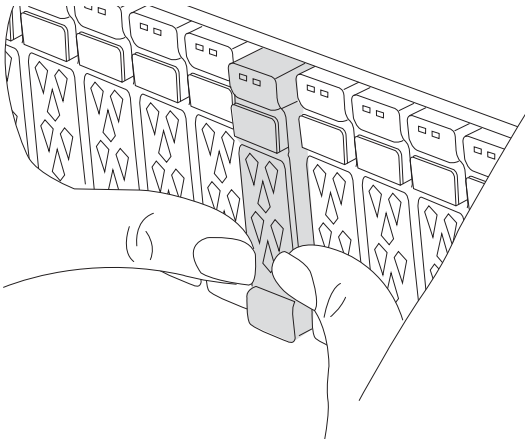

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

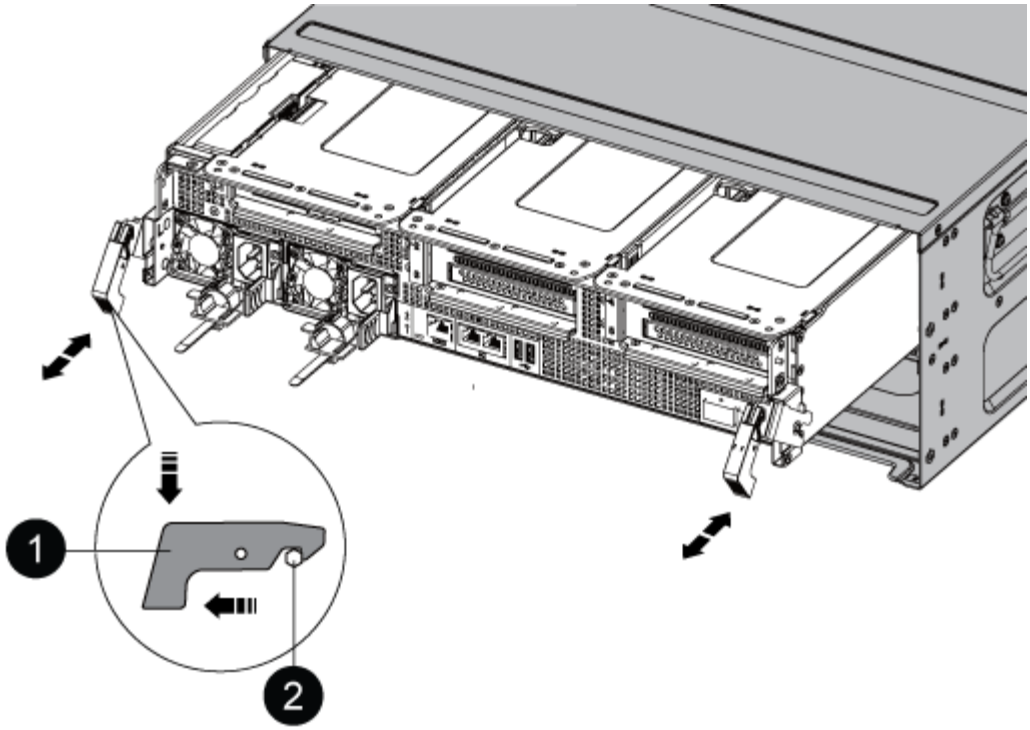

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

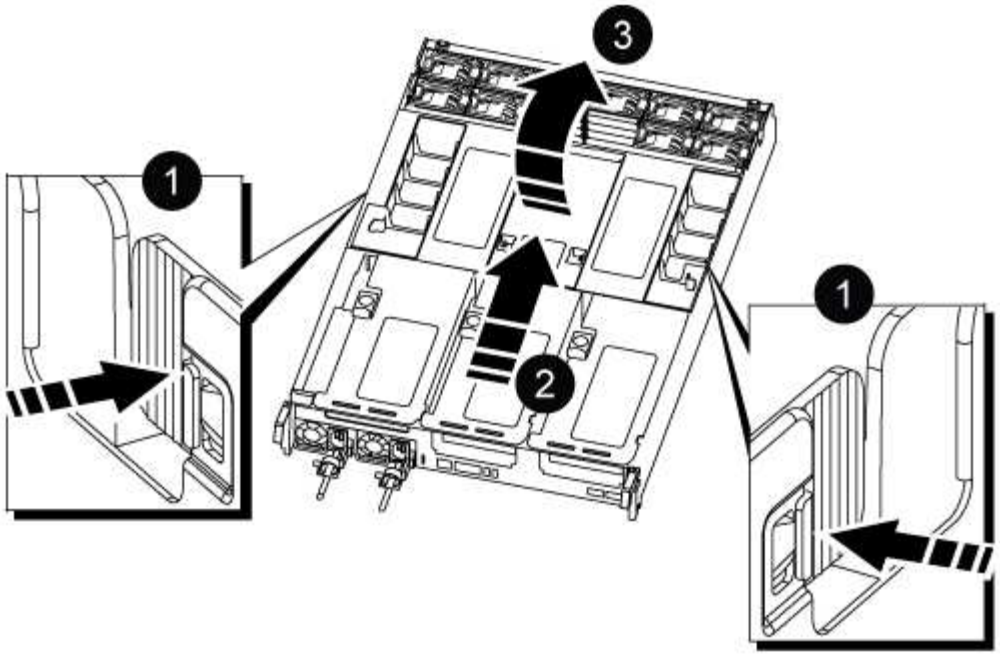

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace a DIMM

To replace a DIMM, you must locate it in the controller module using the DIMM map label on top of the air duct and then replace it following the specific sequence of steps.

1. When removing a DIMM, unlock the locking latch on the applicable riser, and then remove the riser.

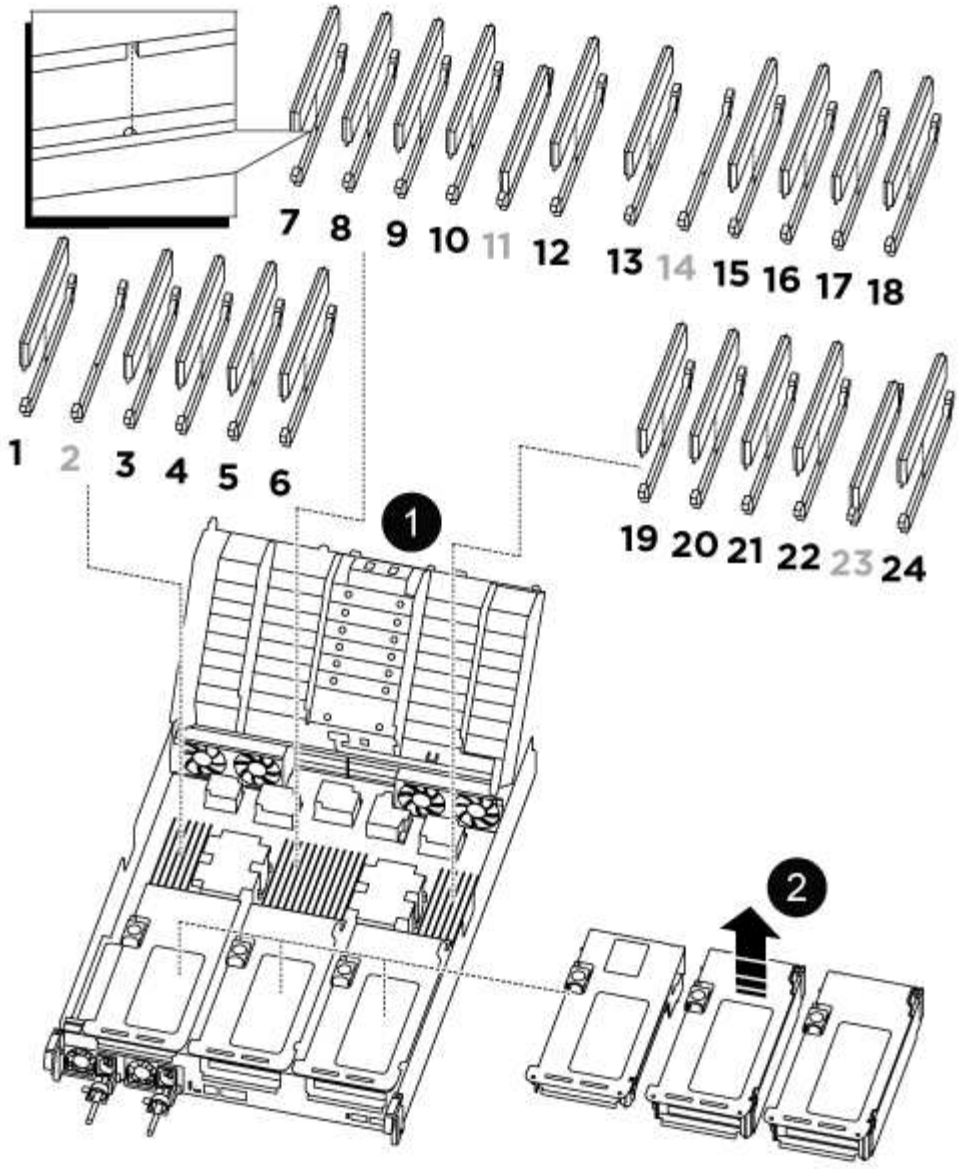

1	Air duct cover
2	Riser 1 and DIMM bank 1, and 3-6
Riser 2 and DIMM bank 7-10, 12-13, and 15-18	Riser 3 and DIMM 19 -22 and 24

Note: Slot 2 and 14 are left empty. Do not attempt to install DIMMs into these slots.

2. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

4. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

5. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Reinstall any risers that you removed from the controller module.
8. Close the air duct.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

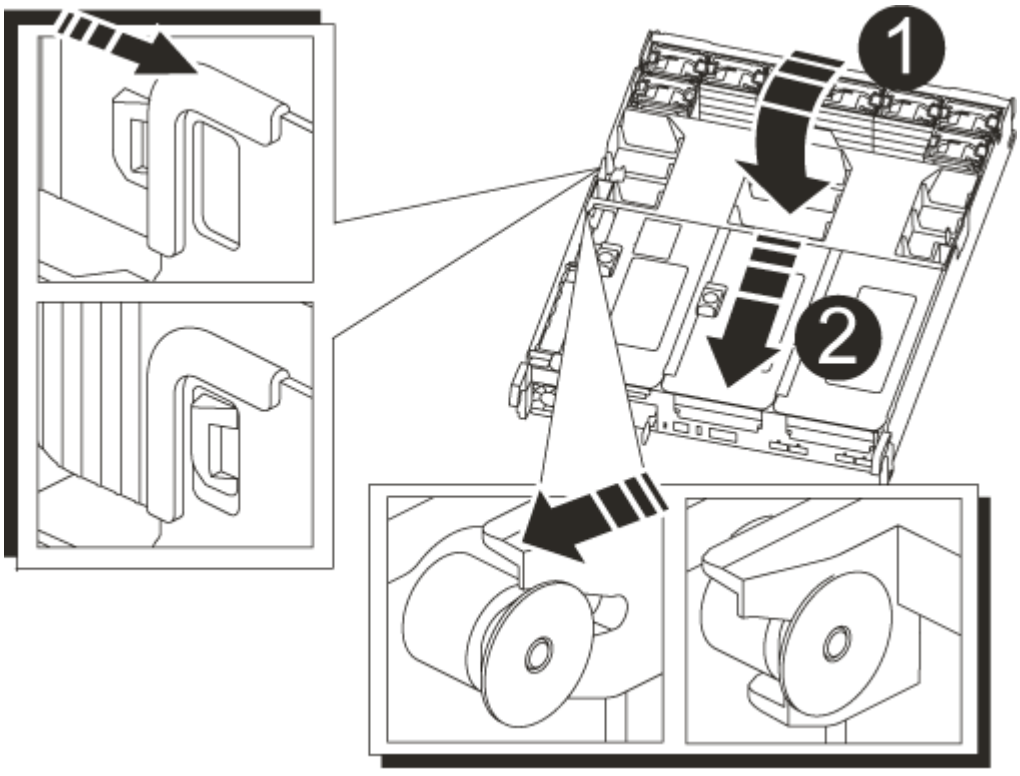

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD drive - AFF C800

You can replace a failed SSD drive nondisruptively while I/O is in progress.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed disk drive by running the `storage disk show -broken` command from the system

console.

The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

It can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the disk depends on how the disk drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

When replacing several disk drives, you must wait 70 seconds between the removal of each failed disk drive and the insertion of the replacement disk drive to allow the storage system to recognize the existence of each new disk.

You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.

2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

7. Verify that the drive's activity (green) LED is illuminated.

   When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      ```
      storage disk show -container-type unassigned
      ```

      You can enter the command on either controller module.

   b. Assign each drive:

      ```
      storage disk assign -disk disk_name -owner node_name
      ```

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      ```
      storage disk option modify -node node_name -autoassign on
      ```

      You must reenable automatic drive assignment on both controller modules.

10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Replace a fan - AFF C800

To replace a fan, remove the failed fan module and replace it with a new fan module.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a fan module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

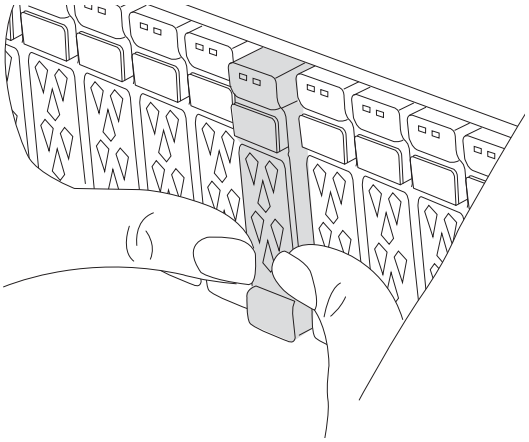

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

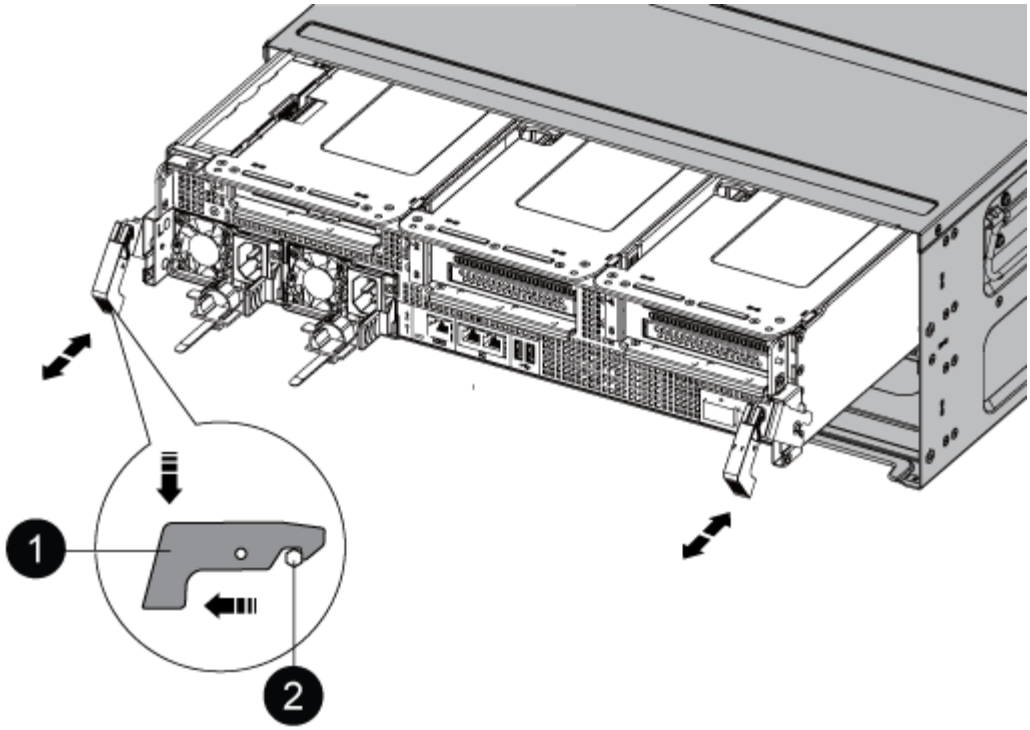

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Set the controller module aside in a safe place.

Step 3: Replace a fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
2. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

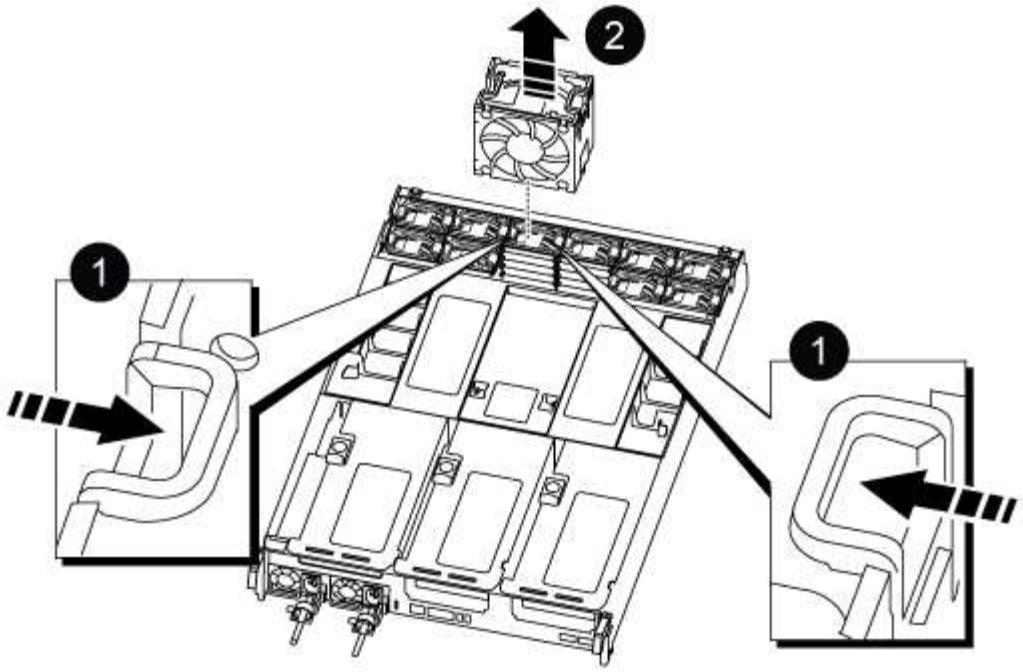

1	Fan locking tabs
2	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module until the locking latches click into place.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.
3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -controller local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an NVDIMM - AFF C800

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

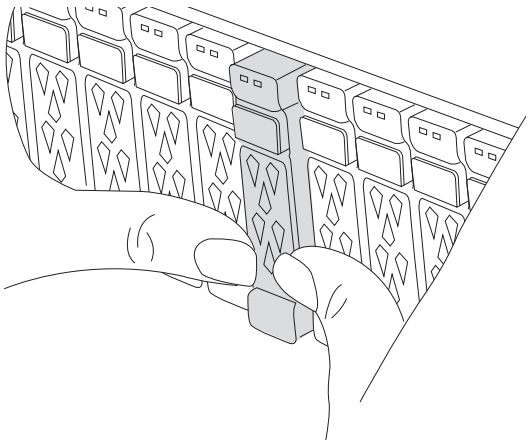

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.

5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

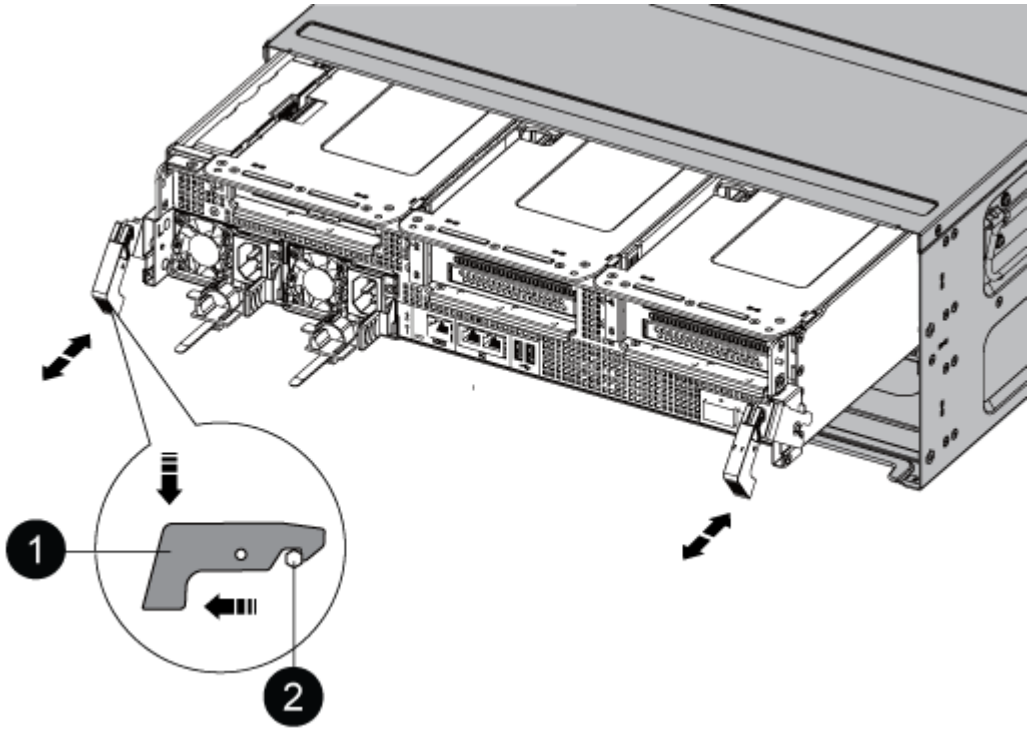

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:

   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

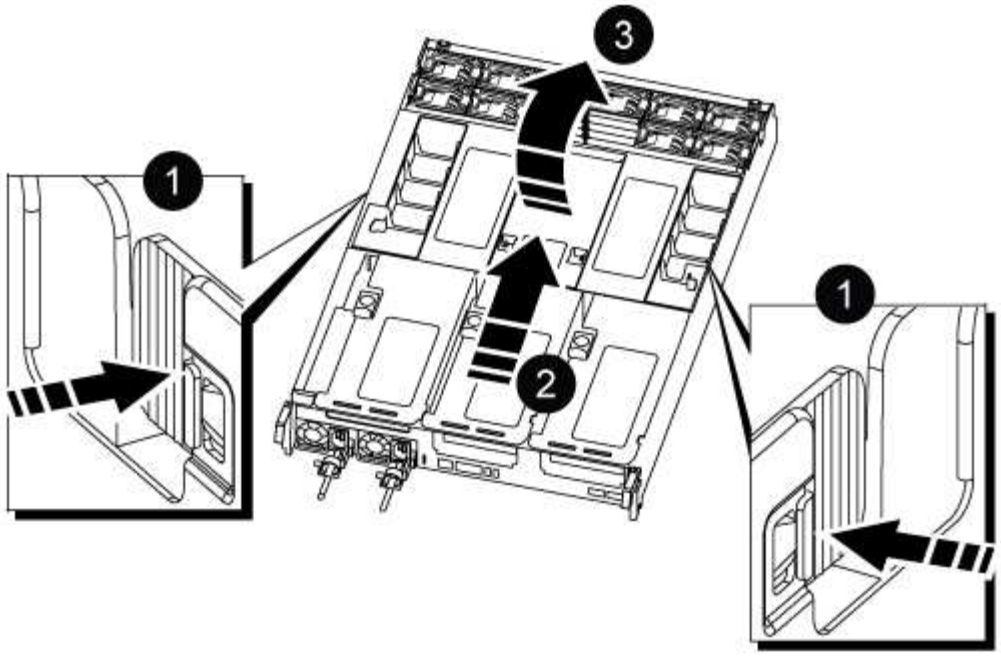

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the NVDIMM map label on top of the air duct, and then replace it following the specific sequence of steps.

1. If you are removing or moving an NVDIMM, unlock the locking latch on the riser, and then remove the applicable riser.

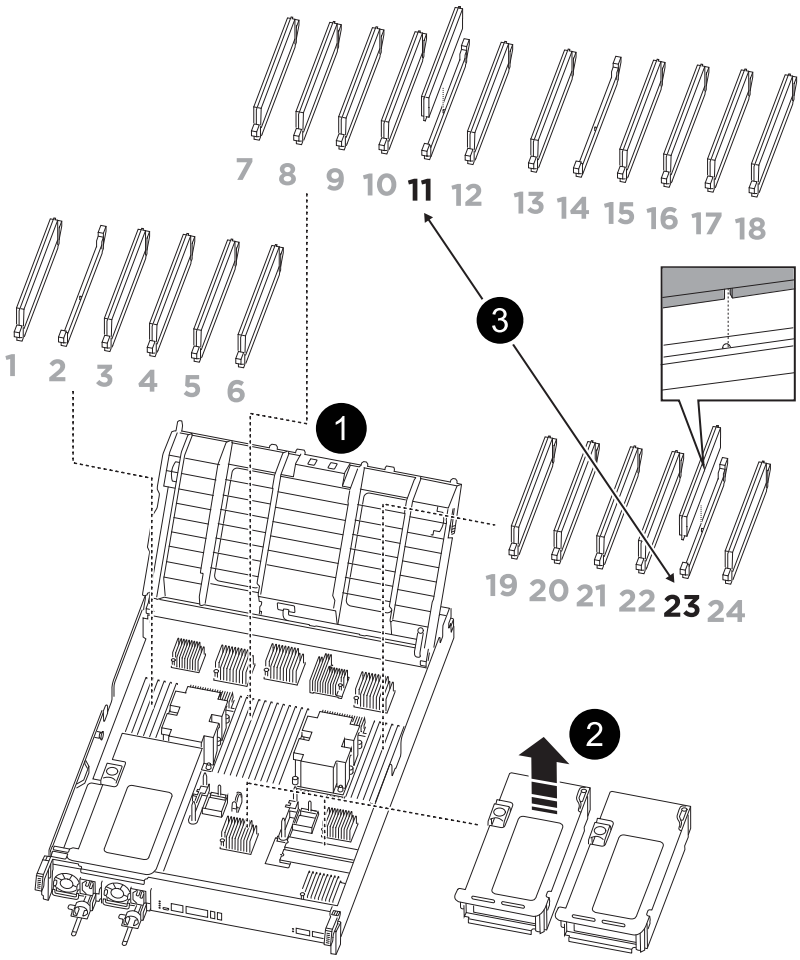

1	Air duct cover
2	Riser 2
3	NVDIMM in slots 11 and 23

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

4. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

5. Locate the slot where you are installing the NVDIMM.

6. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
8. Reinstall any risers that you removed from the controller module.
9. Close the air duct.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

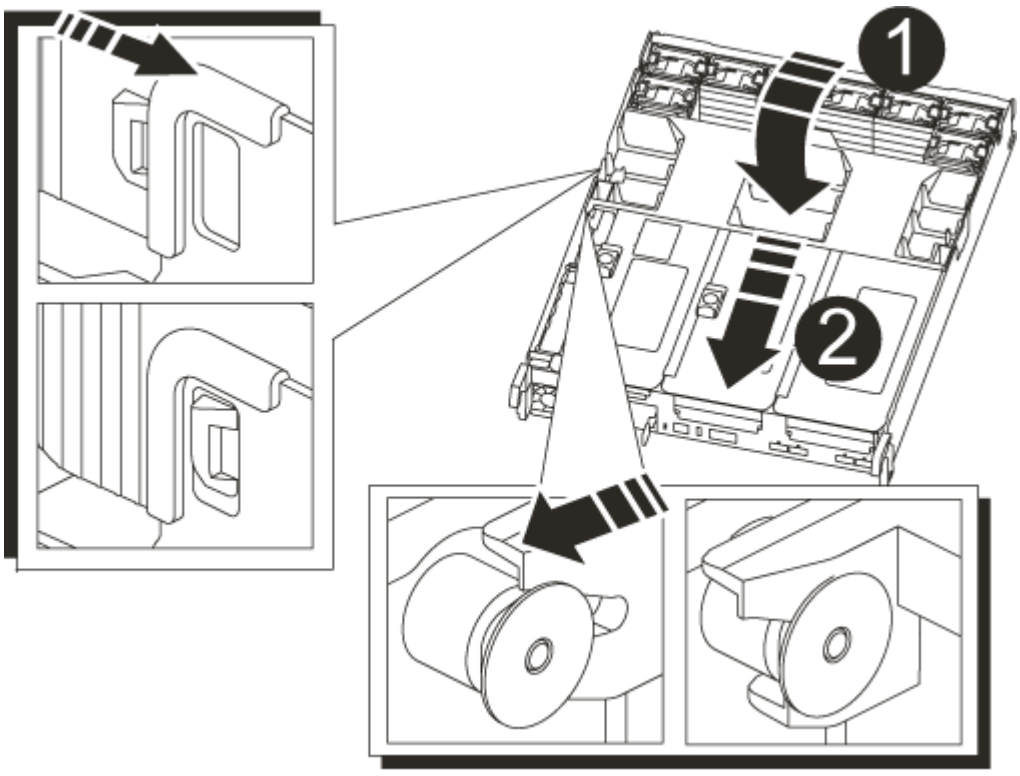

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - AFF C800

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be

resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

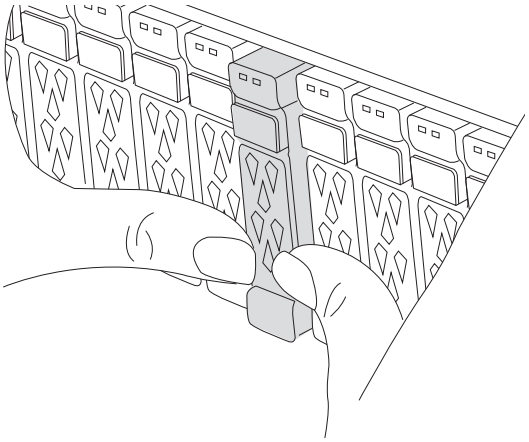

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

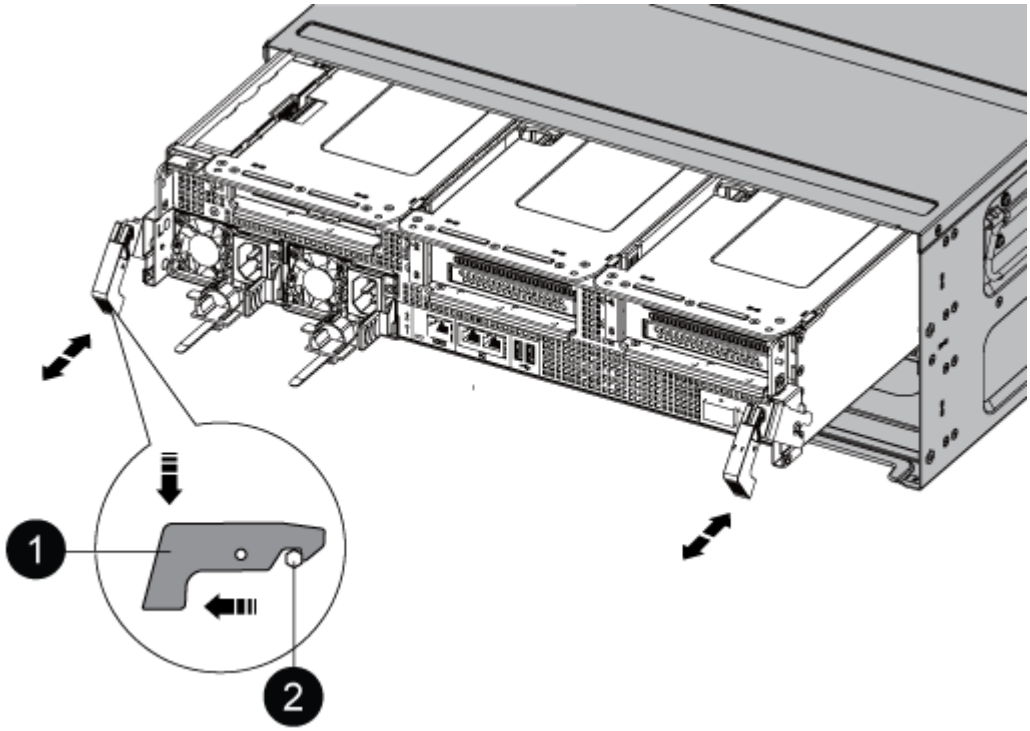

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Set the controller module aside in a safe place.

Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module.

1. Open the air duct cover and locate the NVDIMM battery in the riser.

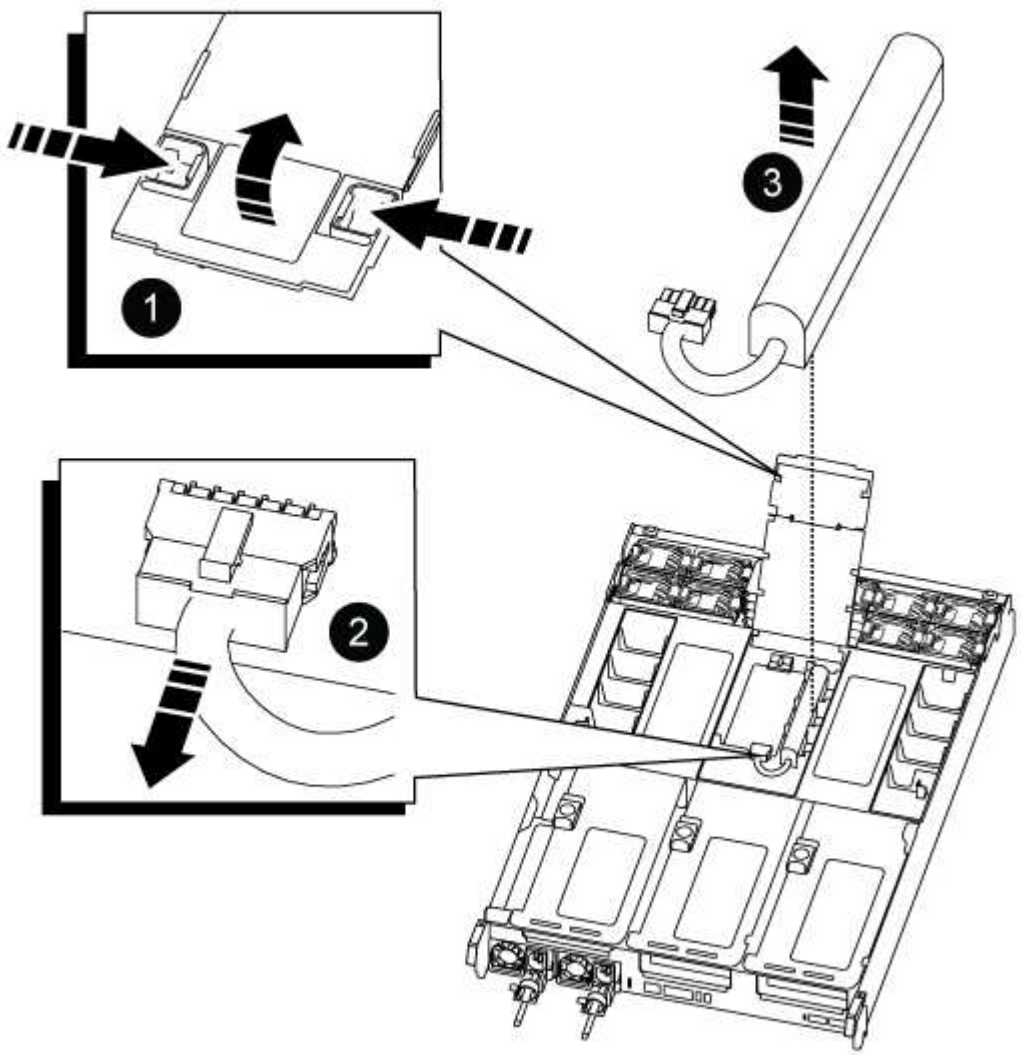

1	Air duct riser

2	NVDIMM battery plug
3	NVDIMM battery pack

Attention: The NVDIMM battery control board LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

1. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
2. Grasp the battery and lift the battery out of the air duct and controller module, and then set it aside.
3. Remove the replacement battery from its package.
4. Install the replacement battery pack in the NVDIMM air duct:
   a. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   b. Plug the battery plug into the riser socket and make sure that the plug locks into place.
5. Close the NVDIMM air duct.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe card - AFF C800

To replace a PCIe card, you must disconnect the cables from the cards, remove the SFP and QSFP modules from the cards before removing the riser, reinstall the riser, and then reinstall the SFP and QSFP modules before cabling the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

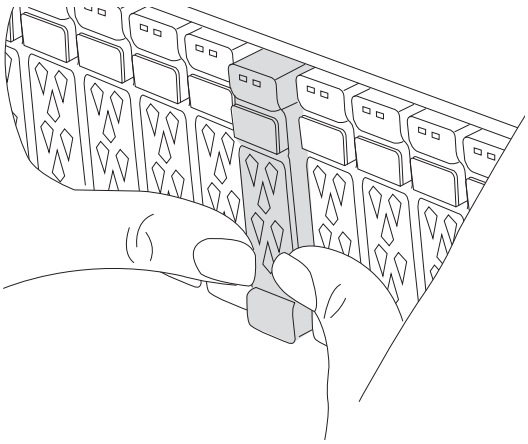

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

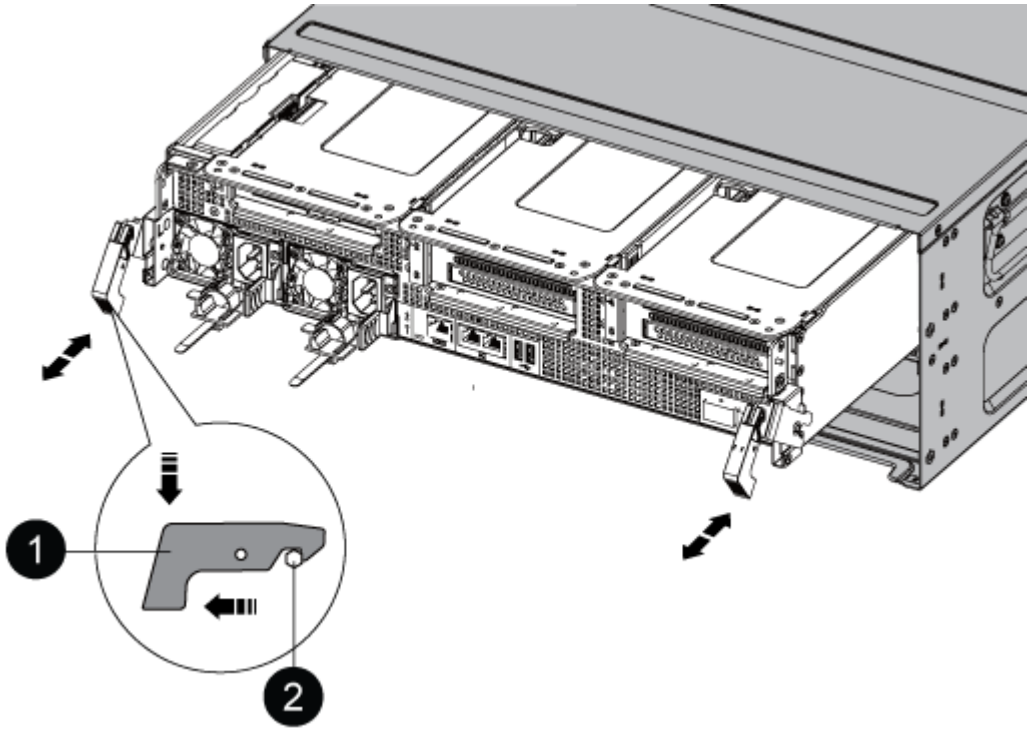

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

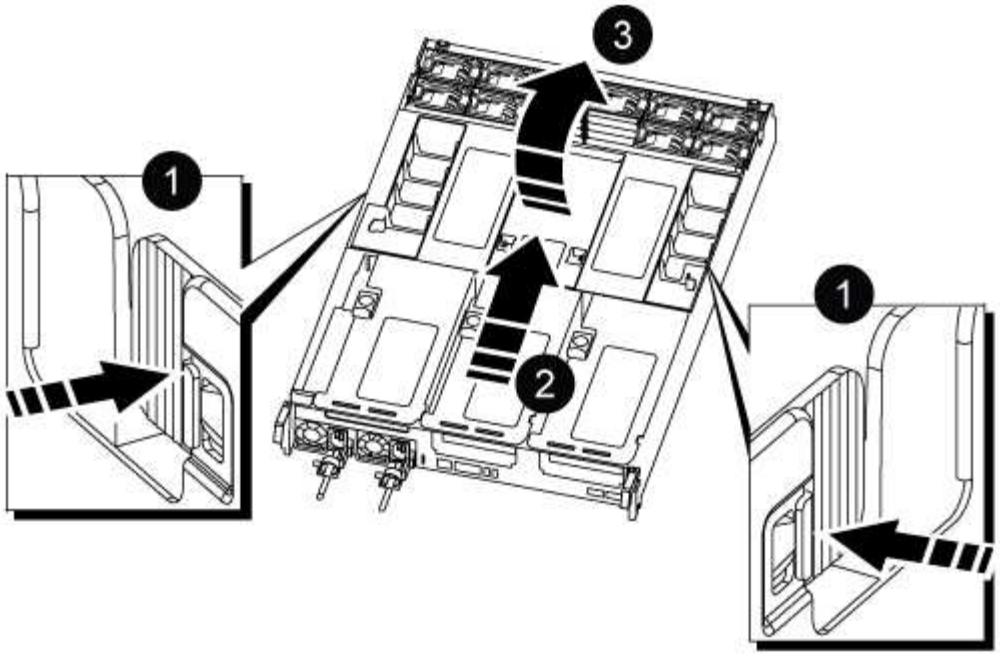

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace a PCIe card

To replace a PCIe card, you must remove the cabling and any QSFPs and SFPs from the ports on the PCIe cards in the target riser, remove the riser from the controller module, remove and replace the PCIe card, reinstall the riser and any QSFPs and SFPs onto the ports, and cable the ports.

1. Determine if the card you are replacing is from Riser 1 or if it is from Riser 2 or 3.
   - If you are replacing the 100GbE PCIe card in Riser 1, use Steps 2 - 3 and Steps 6 - 7.
   - If you are replacing a PCIe card from Riser 2 or 3, use Steps 4 through 7.
2. Remove Riser 1 from the controller module:
   a. Remove the QSFP modules that might be in the PCIe card.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

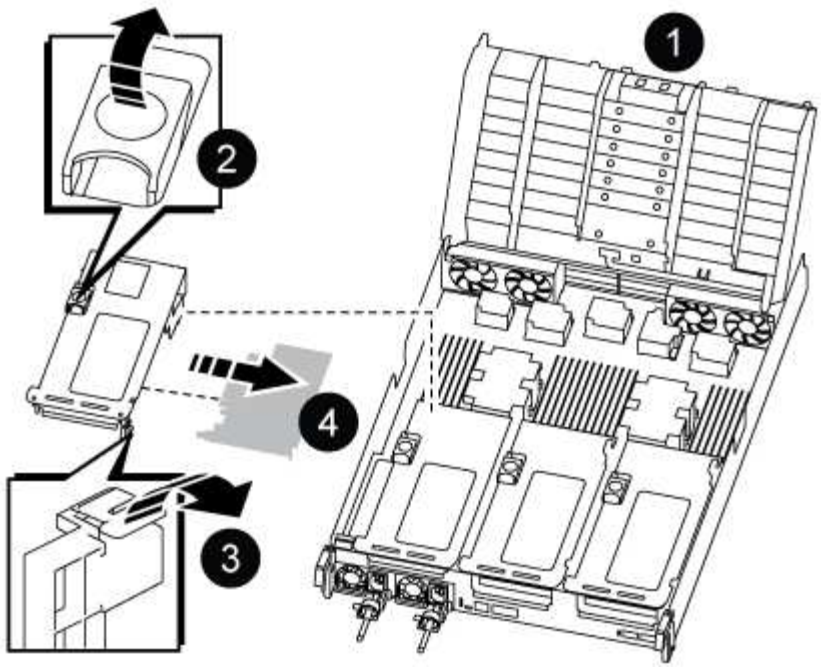

1	Air duct
2	Riser locking latch
3	Card locking bracket
4	Riser 1 (left riser) with 100GbE PCIe card in slot 1.

3. Remove the PCIe card from Riser 1:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. Remove the PCIe card from the riser.
4. Remove the PCIe riser from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

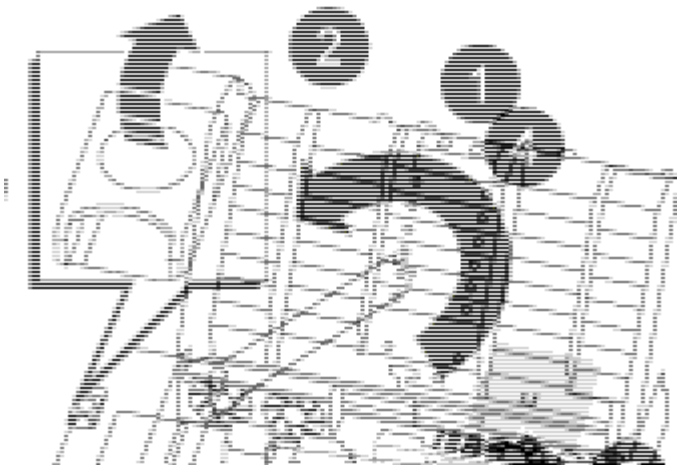

1	Air duct
2	Riser 2 (middle riser) or 3 (right riser) locking latch
3	Card locking bracket
4	Side panel on riser 2 or 3
5	PCIe cards in riser 2 or 3

5. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe cards.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. Swing the side panel off the riser.
   d. Remove the PCIe card from the riser.
6. Install the PCIe card into the same slot in the riser:
   a. Align the card with the card socket in the riser, and then slide it squarely into the socket in the riser.

      Make sure that the card is completely and squarely seated into the riser socket.

   b. For Riser 2 or 3, close the side panel.
   c. Swing the locking latch into place until it clicks into the locked position.
7. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the

controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

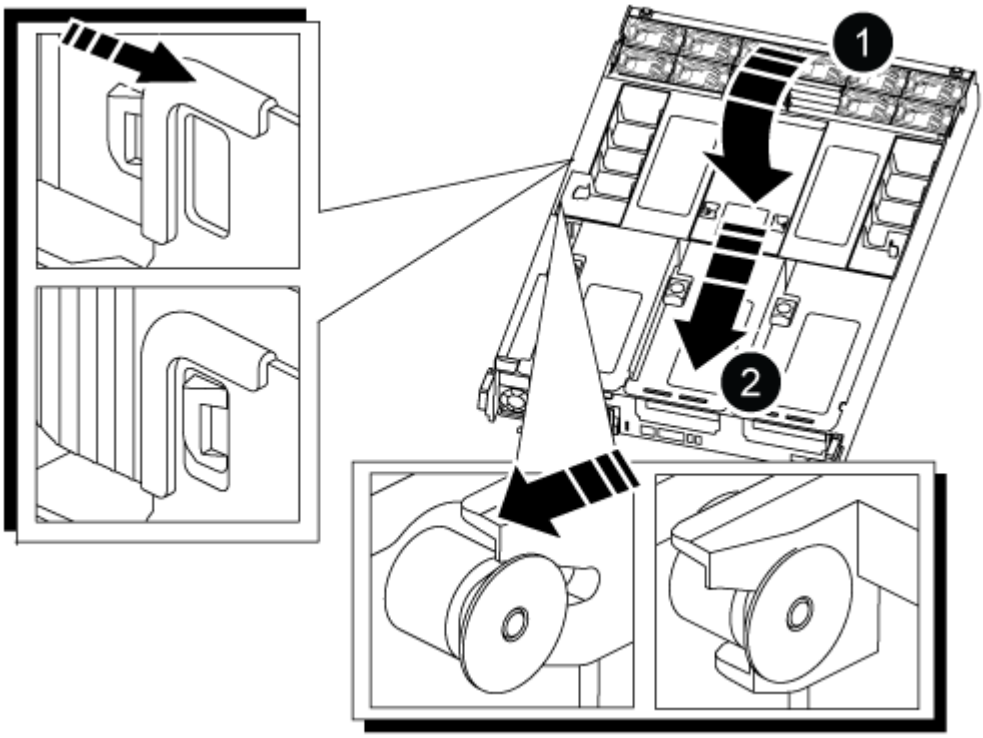

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
5. Plug the system cables and transceiver modules into the controller module and reinstall the cable management device.
6. Plug the power cables into the power supplies and reinstall the power cable retainers.

   The controller module begins to boot as soon as it is connected to power. Be prepared to interrupt the boot process.

   > If your system has DC power supplies, make sure the thumbscrews on the power supply cable are tight.

7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - AFF C800

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

About this task

This procedure is written for replacing one PSU at a time.

> It is a best practice to replace the PSU within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded PSU until the PSU is replaced.

> Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
   b. Unplug the power cable from the power source.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

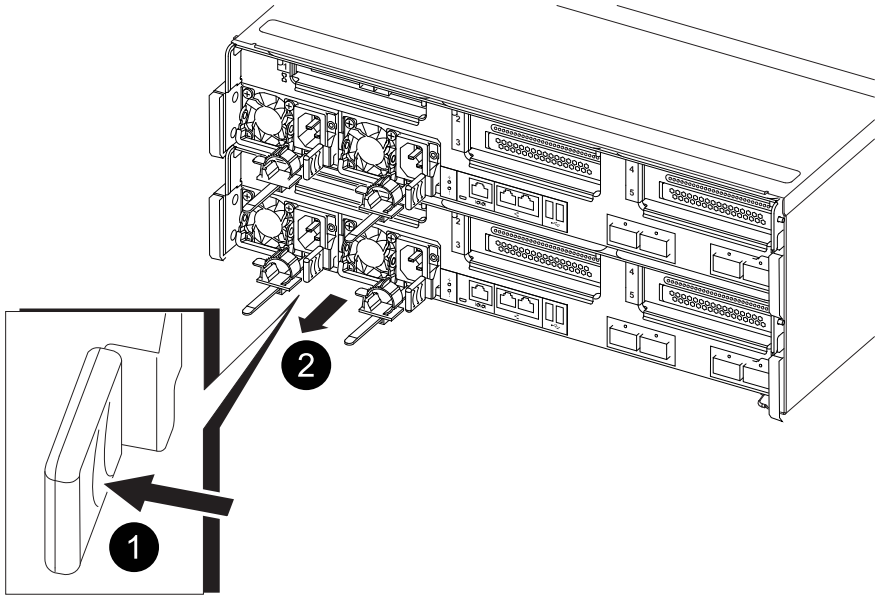

1	Blue PSU locking tab
2	Power supply

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU and the power source.
   b. Secure the power cable to the PSU using the power cable retainer.

   Once power is restored to the PSU, the status LED should be green.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

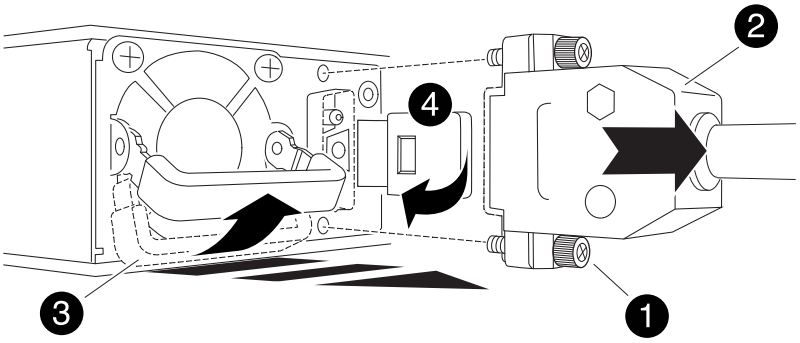

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle

4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF C800

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

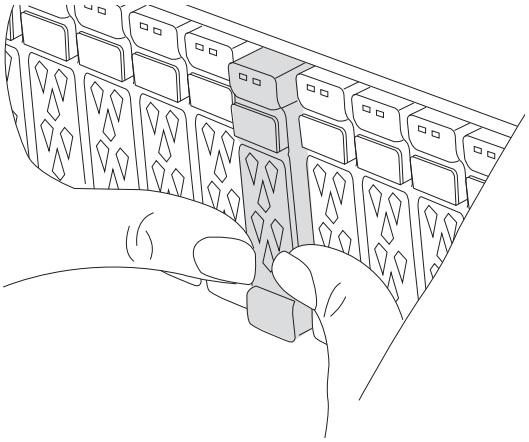

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

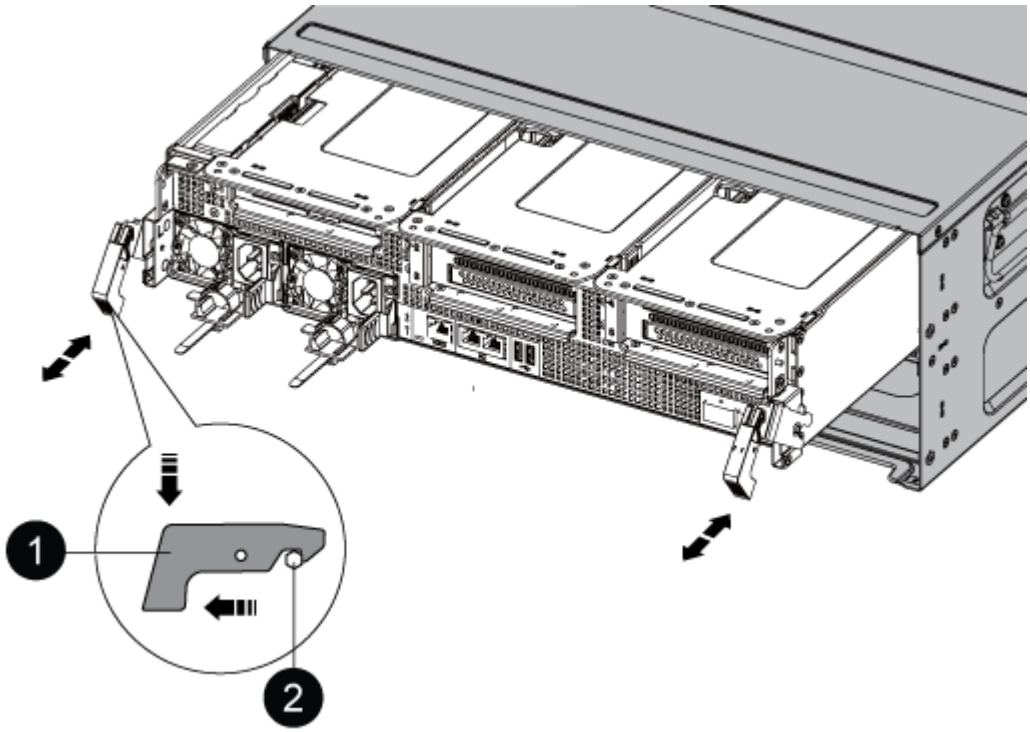

1	Locking latch
2	Locking pin

1. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

2. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

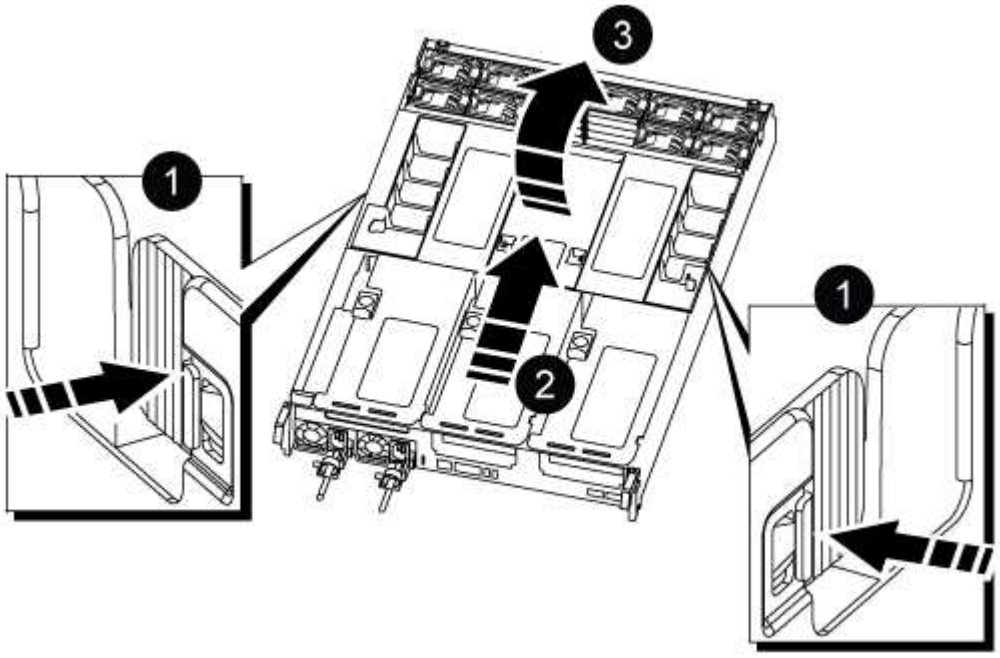

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace the RTC battery

Original controller

1. Remove PCIe riser 2 (middle riser) from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

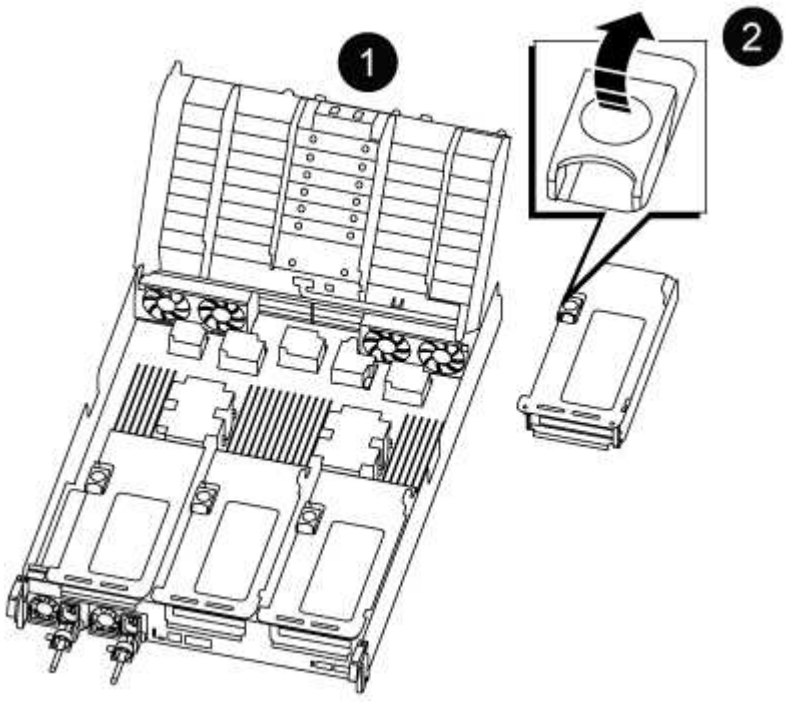

1	Air duct
2	Riser 2 (middle riser) locking latch

2. Locate the RTC battery under Riser 2.

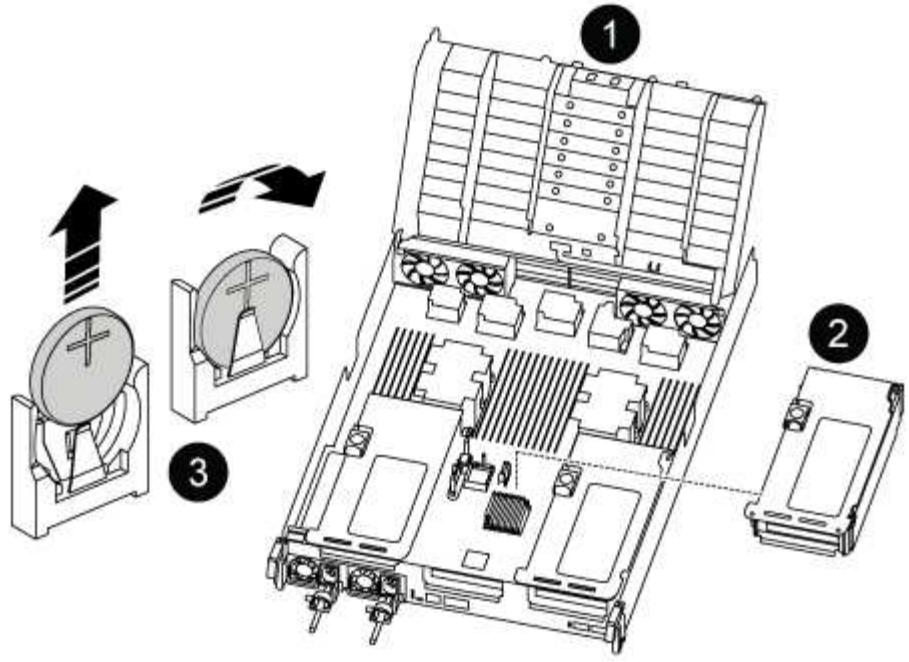

1	Air duct
2	Riser 2
3	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
7. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

VER2 controller

1. Locate the RTC battery near the DIMMs.

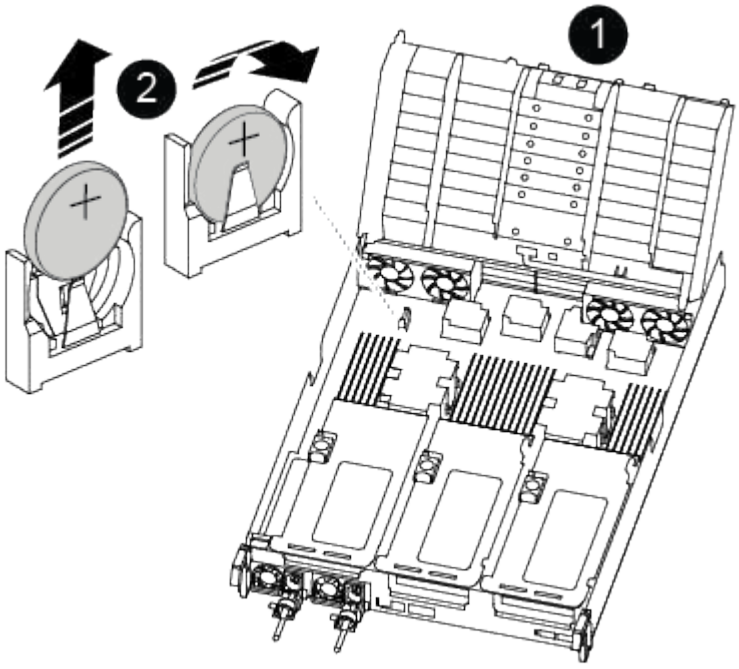

1	Air duct
2	RTC battery and housing

2. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

3. Remove the replacement battery from the antistatic shipping bag.
4. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
5. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Halt the controller at the LOADER prompt.
5. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA systems

ASA A-Series systems

ASA A150 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

Quick guide - ASA A150

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

The Installation and Setup instructions give graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Use the xref:./asa150/AFF A150 System Installation and Setup Instructions

The ASA A150 uses the same installation procedure as the AFF A150 system.

Video steps - ASA A150

Use the following videos to learn how to rack and cable your system and perform initial system configuration.

If you have a MetroCluster configuration, use the MetroCluster documentation.

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

Hardware installation and cabling

The following video shows how to install and cable your system.

Animation - Install and setup of an AFF A150

The ASA A150 uses the same installation procedure as the AFF A150 system.

Detailed guide - ASA A150

Learn how to install your ASA A150 system.

If you have a MetroCluster configuration, use the MetroCluster documentation.

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

Step 1: Prepare for installation

To install your system, create an account on the NetApp Support Site, register your system, and obtain your license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

- Make sure you have access to NetApp Hardware Universe (HWU) for information about site requirements as well as additional information on your configured system.
- Make sure you have access to the Release Notes for your version of ONTAP for more information about this system.
- Contact your network administrator for information about connecting your system to the switches.
- Make sure you have the following items at your site:
  - Rack space for the storage system
  - Phillips #2 screwdriver
  - Additional networking cables to connect your system to your network switch and laptop or console with a Web browser
  - A laptop or console with an RJ-45 connection and access to a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register your system.
4. Download and install Config Advisor on your laptop.
5. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
10 GbE cable (order dependent)	X6566B-05-R6 (112-00297), 0.5m X6566B-2-R6 (112-00299), 2m		Cluster interconnect network
10 GbE cable (order dependent)	Part number X6566B-2-R6 (112-00299), 2m or X6566B-3-R6 (112-00300), 3m X6566B-5-R6 (112-00301), 5m		Data
Optical network cables (order dependent)	X6553-R6 (112-00188), 2m X6536-R6 (112-00090), 5m X6554-R6(112-00189), 15m		FC host network
Cat 6, RJ-45 (order dependent)	Part numbers X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network and Ethernet data
Storage (order dependent)	Part number X66030A (112-00435), 0.5m X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		Storage
Micro-USB console cable	Not applicable		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	Not applicable		Powering up the system

6. Download and complete the Cluster Configuration Worksheet.

Step 2: Install the hardware

Install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

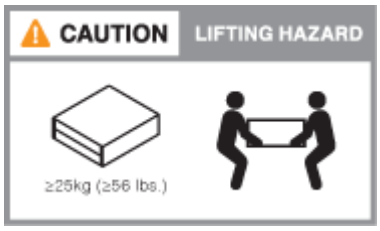

3. Attach cable management devices (as shown).

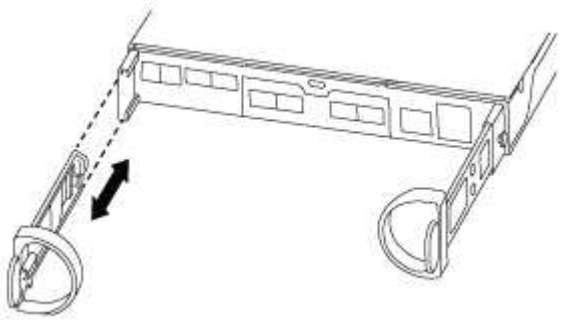

4. Place the bezel on the front of the system.

Step 3: Cable controllers to network

Cable the controllers to your network by using either the two-node switchless cluster method or the cluster interconnect network method.

Management network, UTA2 data network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

Option 1: Two-node switchless cluster

Learn how to cable a two-node switchless cluster.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

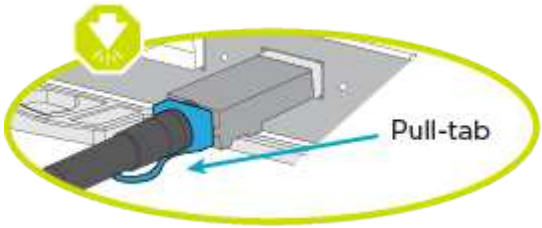

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

You can use either the UTA2 data network ports or the ethernet data network ports to connect the controllers to your host network. Refer to the following cabling illustrations when cabling between the controllers and the switches.

UTA2 data network configurations

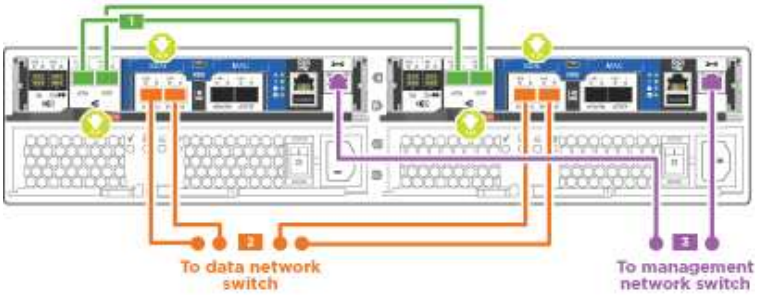

Ethernet network configurations

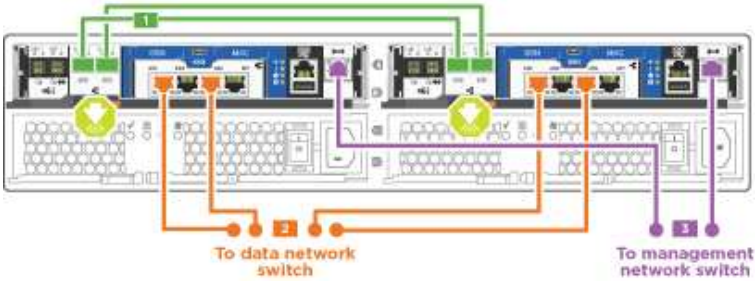

Perform the following steps on each controller module.

Steps

1. Cable the cluster interconnect ports e0a to e0a and e0b to e0b with the cluster interconnect cable.

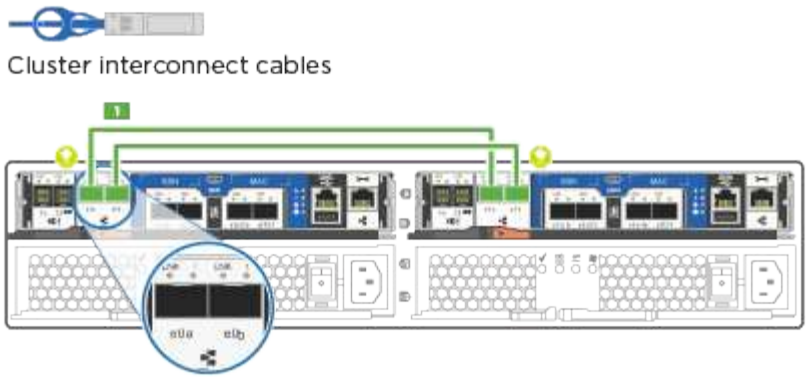

2. Do one of the following:

UTA2 data network configurations

Use one of the following cable types to cable the UTA2 data ports to your host network.

- For an FC host, use 0c and 0d **or** 0e and 0f.
- For an 10GbE system, use e0c and e0d **or** e0e and e0f.

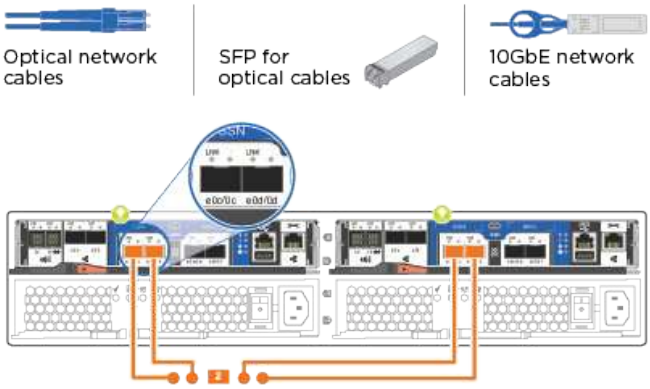

You can connect one port pair as CNA and one port pair as FC, or you can connect both port pairs as CNA or both port pairs as FC.

Ethernet network configurations

Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network.
in the following illustration.

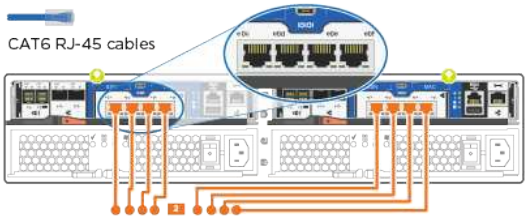

3. Cable the e0M ports to the management network switches with the RJ45 cables.

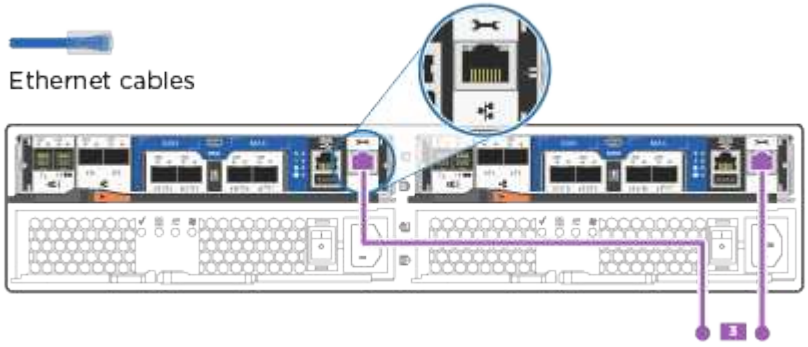

DO NOT plug in the power cords at this point.

Option 2: Switched cluster

Learn how to cable a switched cluster.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

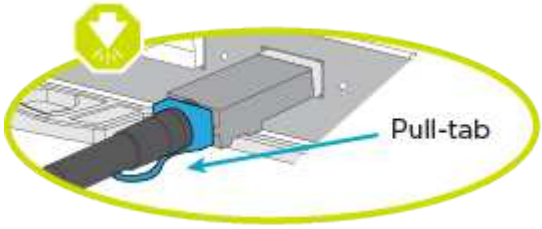

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

You can use either the UTA2 data network ports or the ethernet data network ports to connect the controllers to your host network. Refer to the following cabling illustrations when cabling between the controllers and the switches.

Unified network cabling

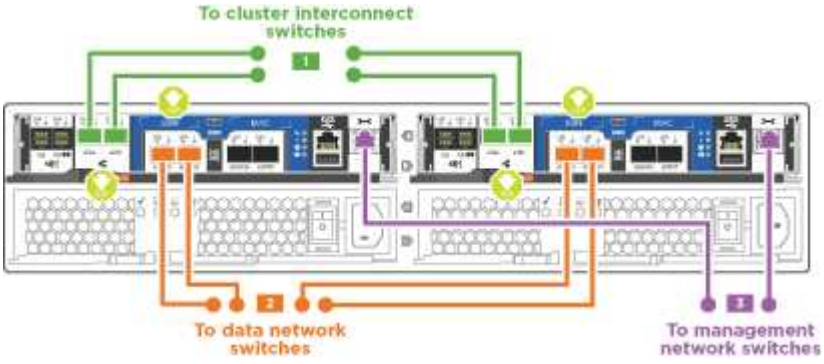

Ethernet network cabling

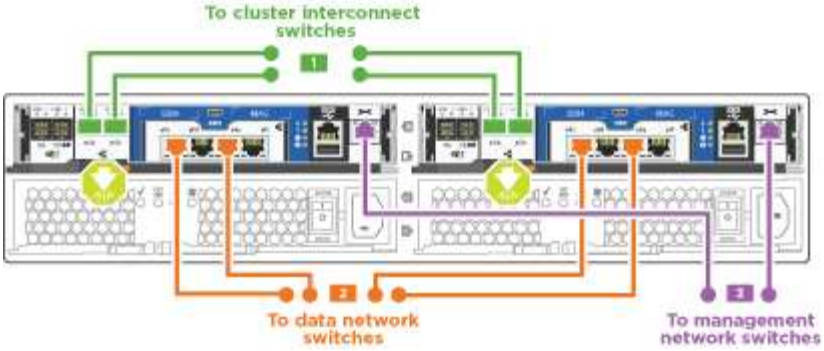

Perform the following steps on each controller module.

Steps

1. For each controller module, cable e0a and e0b to the cluster interconnect switches with the cluster interconnect cable.

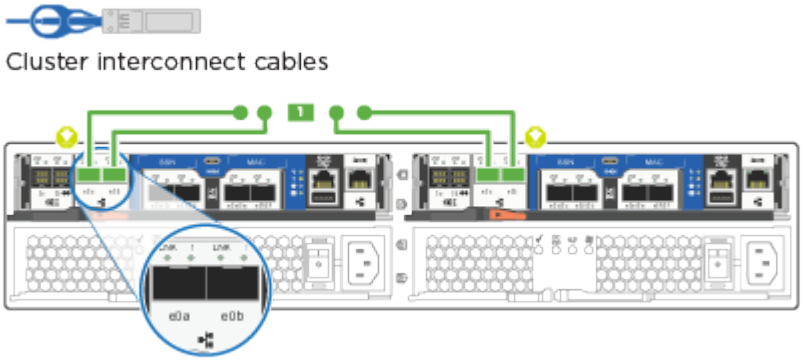

2. Do one of the following:

UTA2 data network configurations

Use one of the following cable types to cable the UTA2 data ports to your host network.

- For an FC host, use 0c and 0d **or** 0e and 0f.
- For an 10GbE system, use e0c and e0d **or** e0e and e0f.

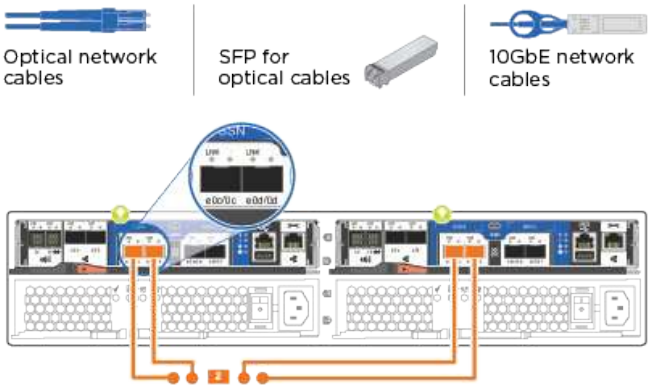

You can connect one port pair as CNA and one port pair as FC, or you can connect both port pairs as CNA or both port pairs as FC.

Ethernet network configurations

Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network.

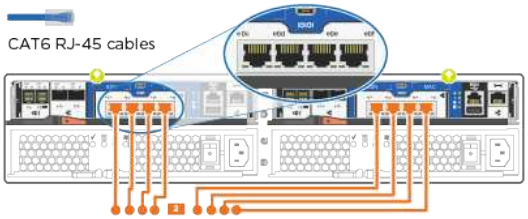

3. Cable the e0M ports to the management network switches with the RJ45 cables.

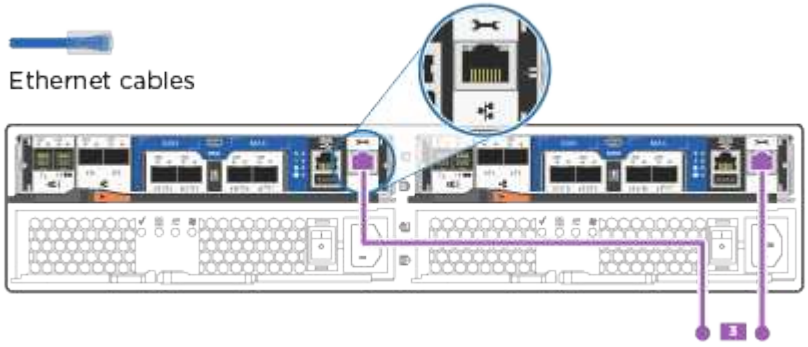

DO NOT plug in the power cords at this point.

Step 4: Cable controllers to drive shelves

Cable the controllers to your shelves using the onboard storage ports. NetApp recommends MP-HA cabling for systems with external storage.

About this task

If you have a SAS tape drive, you can use single-path cabling. If you have no external shelves, MP-HA cabling to internal drives is optional (not shown) if the SAS cables are ordered with the system.

You must cable the shelf-to-shelf connections, and then cable both controllers to the drive shelves.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

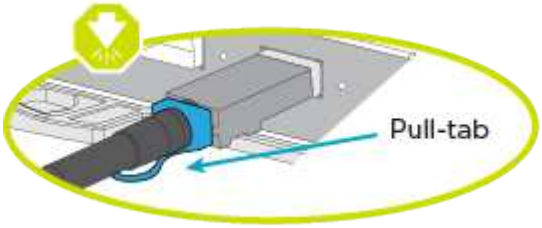

Steps

1. Cable the HA pair with external drive shelves.

   The following example shows cabling for DS224C drive shelves. The cabling is similar with other supported drive shelves.

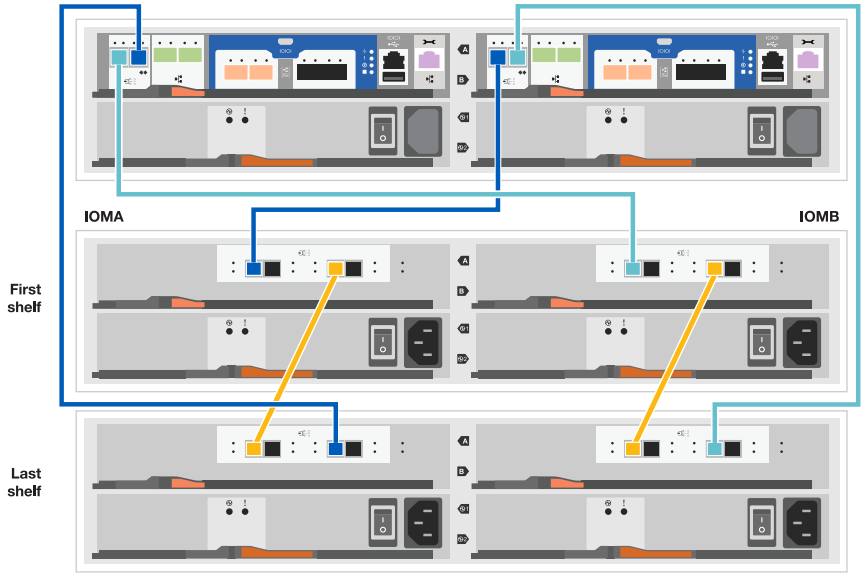

2. Cable the shelf-to-shelf ports.
   - Port 3 on IOM A to port 1 on the IOM A on the shelf directly below.
   - Port 3 on IOM B to port 1 on the IOM B on the shelf directly below.

 mini-SAS HD to mini-SAS HD cables

3. Connect each node to IOM A in the stack.
   - Controller 1 port 0b to IOM A port 3 on last drive shelf in the stack.
   - Controller 2 port 0a to IOM A port 1 on the first drive shelf in the stack.

 mini-SAS HD to mini-SAS HD cables

4. Connect each node to IOM B in the stack
   - Controller 1 port 0a to IOM B port 1 on first drive shelf in the stack.
   - Controller 2 port 0b to IOM B port 3 on the last drive shelf in the stack.

 mini-SAS HD to mini-SAS HD cables

For additional cabling information, see Install and cable shelves for a new system installation - shelves with IOM12/IOM12B modules.

Step 5: Complete system setup

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: If network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Use the following animation to set one or more drive shelf IDs

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Turn on the power switches to both nodes.

Initial booting may take up to eight minutes.

4. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

5. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

6. Select an ONTAP icon listed to discover:

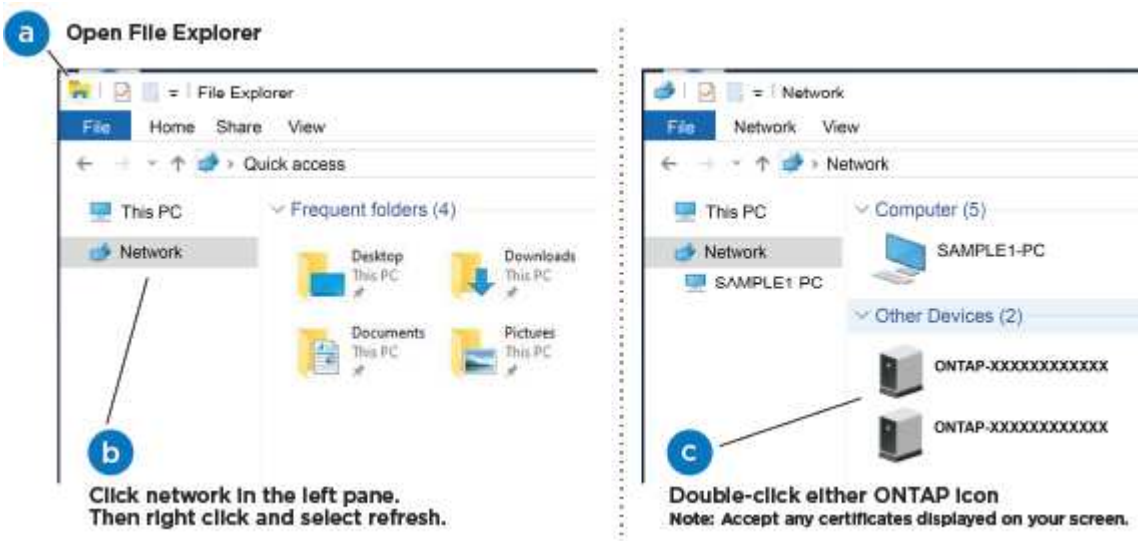

a. Open File Explorer.

b. Click **Network** in the left pane and right-click and select **refresh**.

c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

XXXXX is the system serial number for the target node.

System Manager opens.

7. Configure the system using the data you collected in the ONTAP Configuration Guide.
8. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create and account.
   b. Register your system.
   c. Download Active IQ Config Advisor.
9. Verify the health of your system by running Config Advisor.
10. After you have completed the initial configuration, go to the ONTAP documentation site for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console.
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for instructions on how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your system.

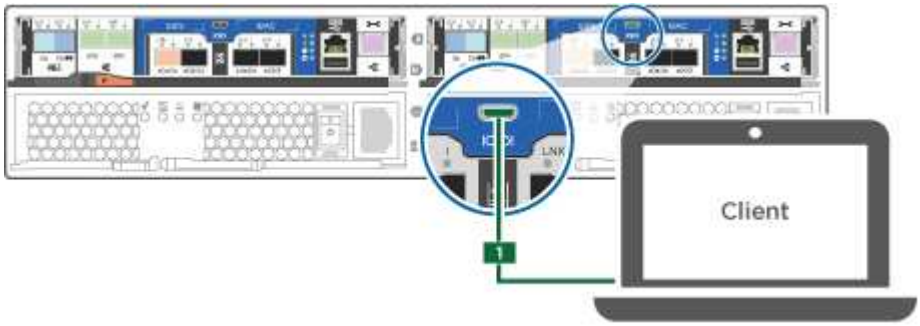

   c. Connect the laptop or console to the switch on the management subnet.

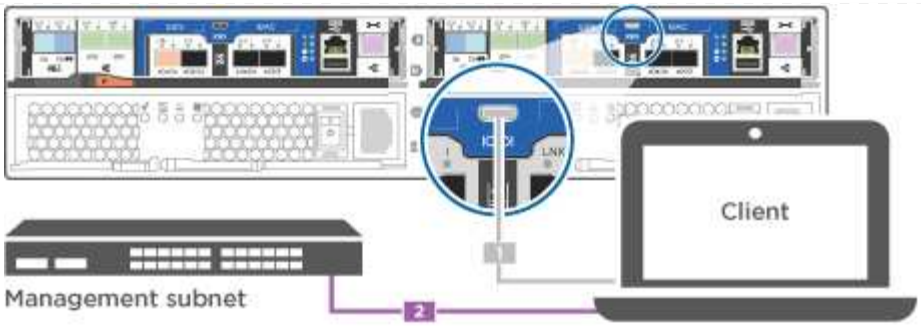

d. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Use the following animation to set one or more drive shelf IDs:

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
4. Turn on the power switches to both nodes.

Initial booting may take up to eight minutes.

5. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

6. Using System Manager on your laptop or console, configure your cluster.
   a. Point your browser to the node management IP address.

The format for the address is https://x.x.x.x.

b. Configure the system using the data you collected in the ONTAP Configuration Guide.

7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create and account.
   b. Register your system.
   c. Download Active IQ Config Advisor.
8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to the ONTAP documentation site for information about configuring additional features in ONTAP.

Maintain

Maintain ASA A150 hardware

For the ASA A150 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

NVEM Battery

A battery is included with a controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

////July 2025: ontap-systems 370: deleted caching module since not supported for this platform.

Boot media

Overview of boot media replacement - ASA A150

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct node:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy node* is the HA partner of the impaired node.

Check encryption key support and status - ASA A150

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select

one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the impaired controller - ASA A150

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller displays…	Then…
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - ASA A150

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

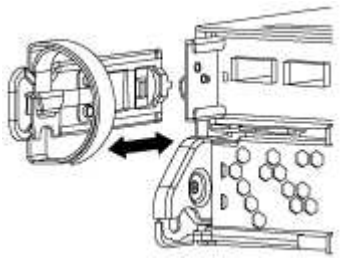

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

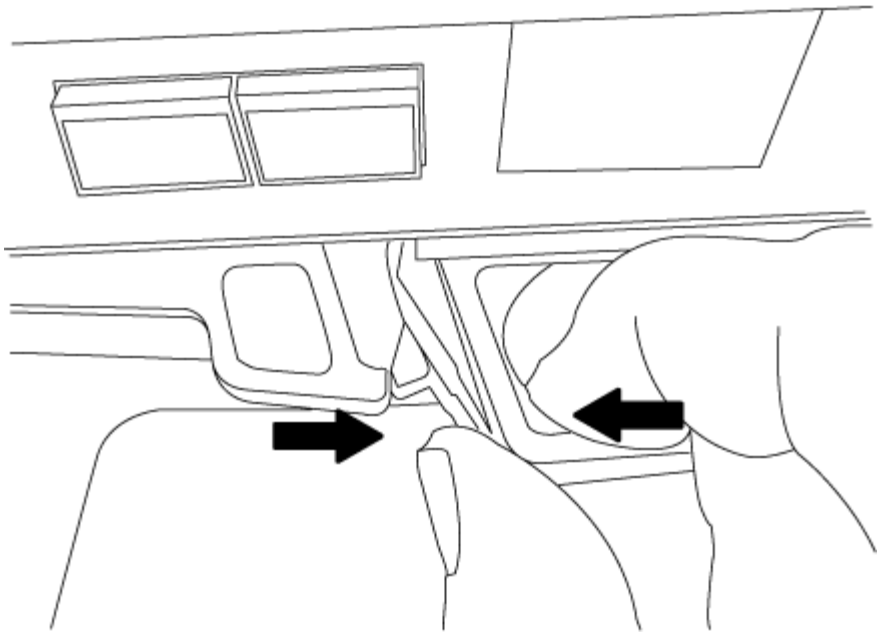

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

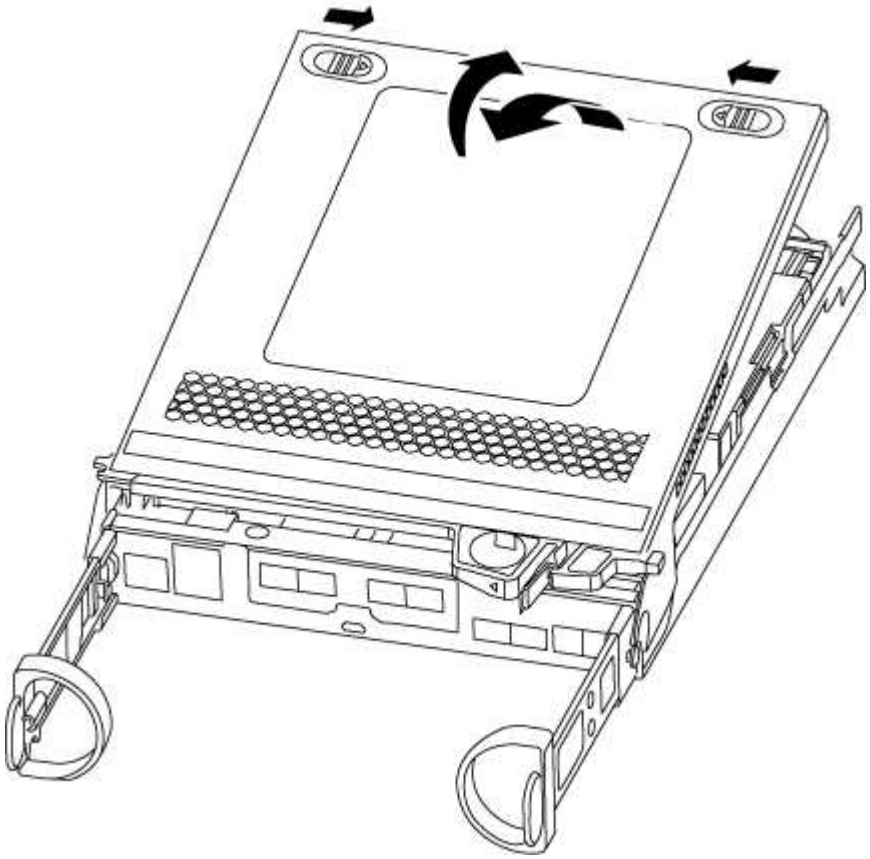

Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the boot media using the following illustration or the FRU map on the controller module:

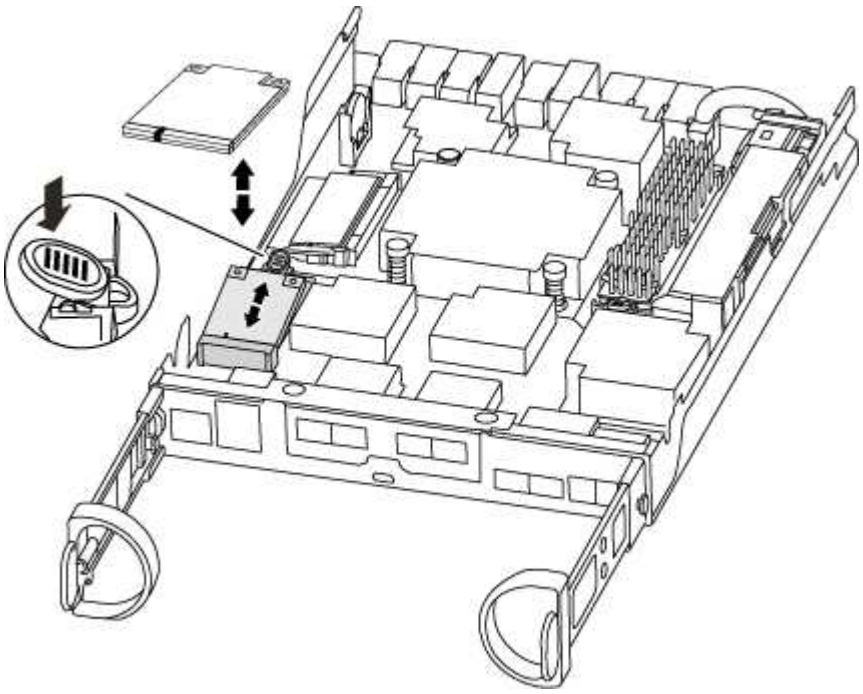

3. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

4. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
5. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

6. Push the boot media down to engage the locking button on the boot media housing.
7. Close the controller module cover.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

   The controller begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

6. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

   The system begins to boot and stops at the LOADER prompt.

7. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

     The target port you configure is the target port you use to communicate with the impaired controller from the healthy controller during var file system restore with a network connection. You can also use the e0M port in this command.

   - If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
     - filer_addr is the IP address of the storage system.
     - netmask is the network mask of the management network that is connected to the HA partner.
     - gateway is the gateway for the network.
     - dns_addr is the IP address of a name server on your network.
     - dns_domain is the Domain Name System (DNS) domain name.

       If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

Boot the recovery image - ASA A150

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:

- If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed
- If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.
8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore OKM, NSE, and NVE - ASA A150

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.

- Backup information for the Onboard Key Manager.

- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option

ONTAP 9.8 or later	Select option 10.
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`(10) Set Onboard Key Manager recovery secrets.`
	`(11) Configure node for external key management.`
	`Selection (1-11)? 10`

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - ASA A150

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - ASA A150

To replace the chassis, you must move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving all drives and controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-controller cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - ASA A150

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Controller is in a MetroCluster configuration

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the chassis - ASA A150

Move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Move a power supply

Moving out a power supply when replacing a chassis involves turning off, disconnecting, and removing the power supply from the old chassis and installing and connecting it on the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.
4. Use the cam handle to slide the power supply out of the system.

   

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

Step 2: Remove the controller module

Remove the controller module or modules from the old chassis.

1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

2. Remove and set aside the cable management devices from the left and right sides of the controller module.

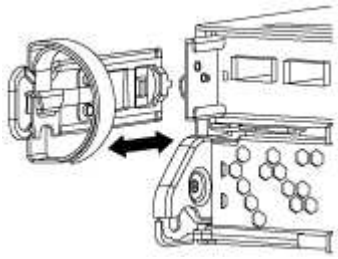

3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

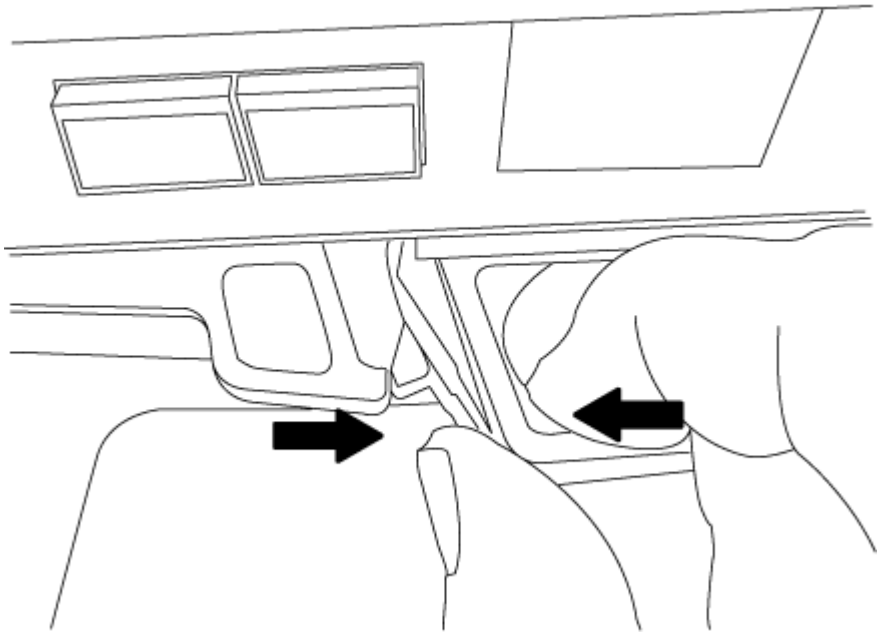

4. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 3: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new

chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.
5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It click when it is secure.
6. Repeat the process for the remaining drives in the system.

Step 4: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller

After you install the controller module and any other components into the new chassis, boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Repeat the preceding steps if there is a second controller to install in the new chassis.
4. Complete the installation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Repeat the preceding steps for the second controller module in the new chassis.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reinstall the blanking panel and then go to the next step.

5. Connect the power supplies to different power sources, and then turn them on.
6. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

> If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - ASA A150

You must verify the HA state of the chassis, switch back aggregates, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis *HA-state*`

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mcc-2n`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. Reboot the system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller replacement - ASA A150

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- This procedure includes steps for automatically or manually reassigning drives to the *replacement* controller, depending on your system’s configuration.

  You should perform the drive reassignment as directed in the procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - ASA A150

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller - ASA A150

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove controller module

To replace the controller module, you must first remove the old controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

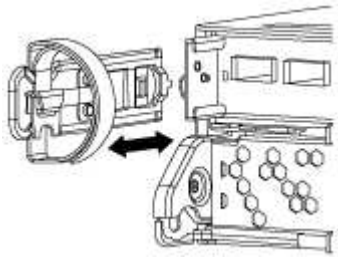

4. If you left the SFP modules in the system after removing the cables, move them to the new controller module.
5. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

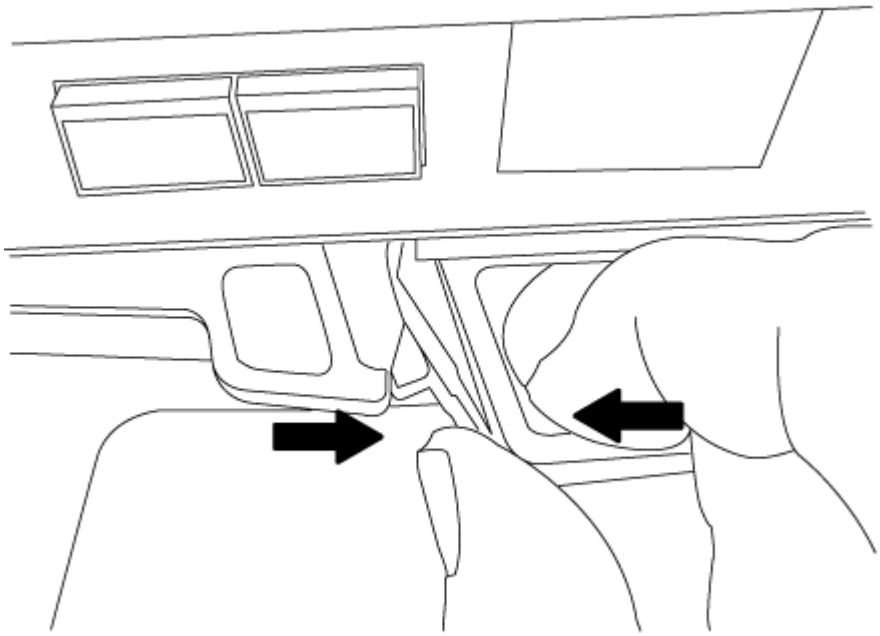

6. Turn the controller module over and place it on a flat, stable surface.
7. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

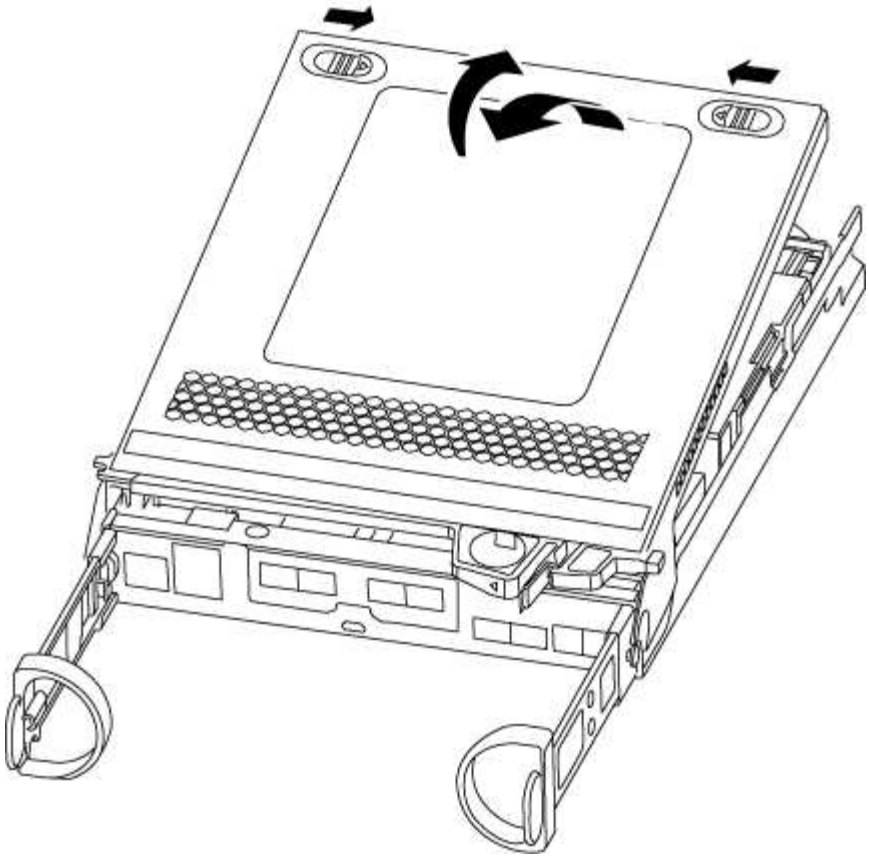

Step 2: Move the NVMEM battery

To move the NVMEM battery from the old controller module to the new controller module, you must perform a specific sequence of steps.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Locate the NVMEM battery in the controller module.

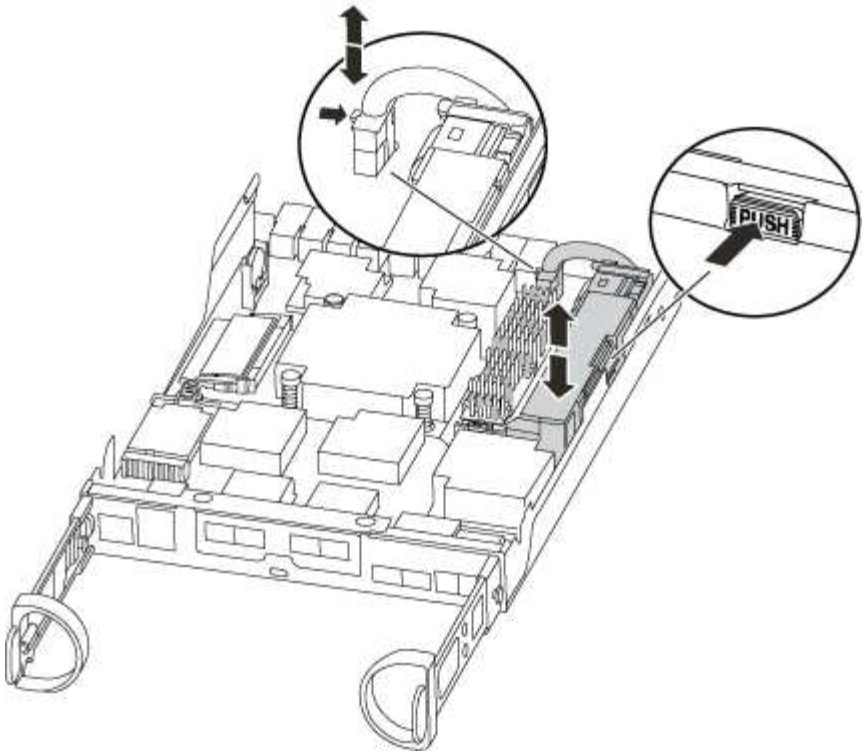

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Loop the battery cable around the cable channel on the side of the battery holder.
7. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
8. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.

Step 3: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller module and insert it in the new controller module.

1. Locate the boot media using the following illustration or the FRU map on the controller module:

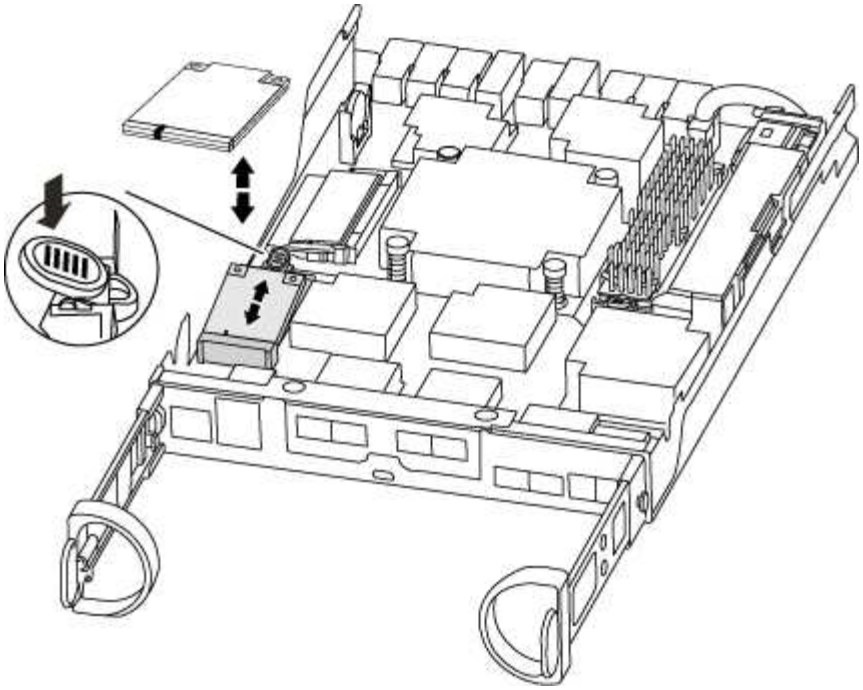

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 4: Move the DIMMs

To move the DIMMs, you must follow the directions to locate and move them from the old controller module into the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

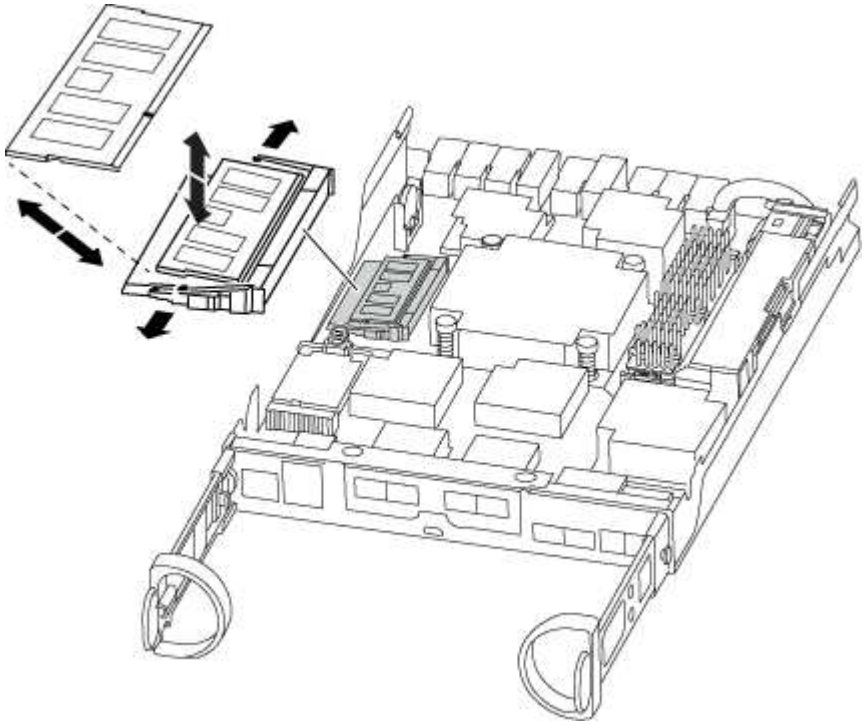

4. Repeat these steps to remove additional DIMMs as needed.
5. Verify that the NVMEM battery is not plugged into the new controller module.
6. Locate the slot where you are installing the DIMM.
7. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Repeat these steps for the remaining DIMMs.
9. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

   Make sure that the plug locks down onto the controller module.

Step 5: Move a caching module, if present

If your AFF A220 or FAS2700 system has a caching module, you need to move the caching module from the old controller module to the replacement controller module. The caching module is referred to as the “M.2 PCIe card” on the controller module label.

You must have the new controller module ready so that you can move the caching module directly from the old controller module to the corresponding slot in the new one. All other components in the storage system must be functioning properly; if not, you must contact technical support.

1. Locate the caching module at the rear of the controller module and remove it.

a. Press the release tab.

b. Remove the heatsink.

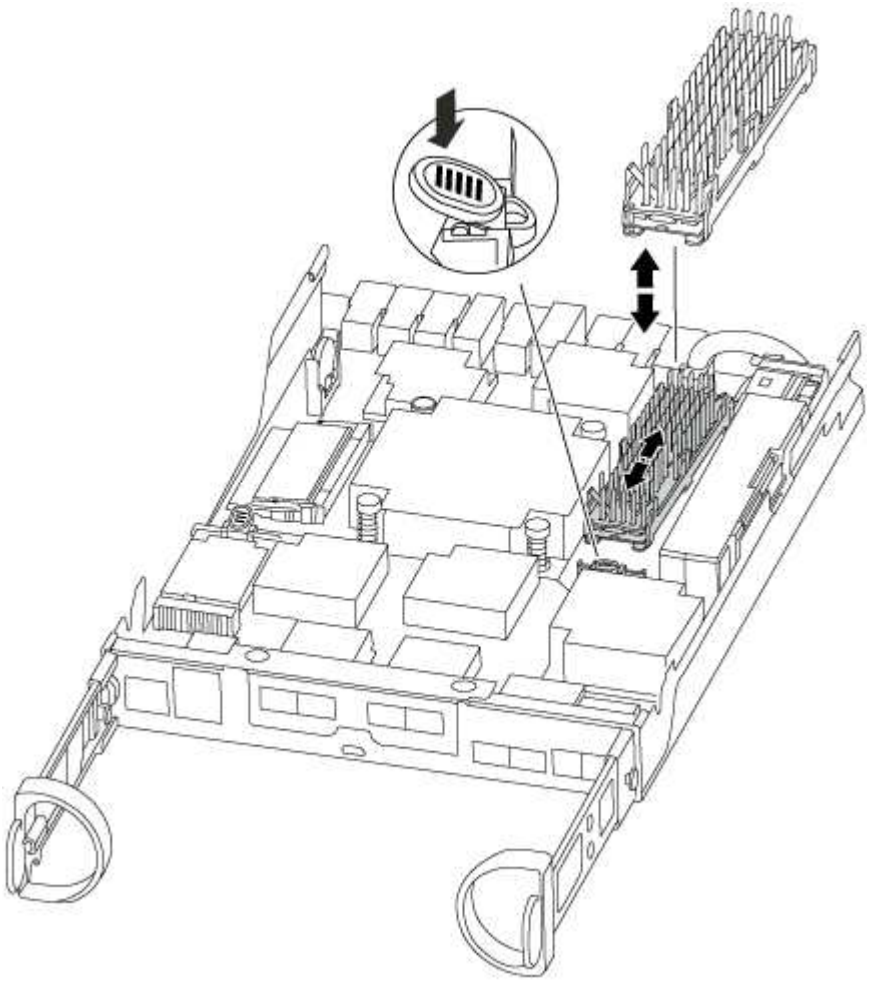

2. Gently pull the caching module straight out of the housing.
3. Move the caching module to the new controller module, and then align the edges of the caching module with the socket housing and gently push it into the socket.
4. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

5. Reseat and push the heatsink down to engage the locking button on the caching module housing.
6. Close the controller module cover, as needed.

Step 6: Install the controller

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Interrupt the boot process **only** after determining the correct timing: You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. Only press `Ctrl-C` when you see the message `Press Ctrl-C for Boot Menu`. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. Select the option to boot to Maintenance mode from the displayed menu.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process. e. Interrupt the boot process **only** after determining the correct timing: You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. Only press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. f. From the boot menu, select the option for Maintenance mode.

Important: During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
  You can safely respond `y` to these prompts.

Restore and verify the system configuration - ASA A150

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

- `ha`
- `mcc`
- `mcc-2n`
- `mccip`
- `non-ha`

b. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - ASA A150

To complete the replacement procedure and restore your system to full operation, you must recable the storage, confirm disk reassignment, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks.

You must use the correct procedure for your configuration:

Controller redundancy	Then use this procedure…
HA pair	Option 1: Verify the system ID change on an HA system
Stand-alone	Option 2: Manually reassign the system ID on a stand-alone system in ONTAP
Two-node MetroCluster configuration	Option 3: Manually reassign the system ID on systems in a two-node MetroCluster configuration

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *`local-node-name`* `partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

- Restore onboard key management encryption keys
- Restore external key management encryption keys

6. Give back the controller:

   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Option 2: Manually reassign the system ID on a stand-alone system in ONTAP

In a stand-alone system, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems that are in a stand-alone configuration.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by pressing Ctrl-C, and then select the option to boot to Maintenance mode from the displayed menu.
2. You must enter `Y` when prompted to override the system ID due to a system ID mismatch.
3. View the system IDs: `disk show -a`
4. You should make a note of the old system ID, which is displayed as part of the disk owner column.

   The following example shows the old system ID of 118073209:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                  POOL   SERIAL NUMBER  HOME
--------    -------------          -----  -------------  -------------
disk_name    system-1  (118073209)  Pool0  J8XJE9LC       system-1
(118073209)
disk_name    system-1  (118073209)  Pool0  J8Y478RC       system-1
(118073209)
.
.
.
```

5. Reassign disk ownership by using the system ID information obtained from the disk show command: `disk reassign -s old system ID disk reassign -s 118073209`
6. Verify that the disks were assigned correctly: `disk show -a`

   The disks belonging to the replacement node should show the new system ID. The following example now show the disks owned by system-1 the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                  POOL   SERIAL NUMBER  HOME
--------    -------------          -----  -------------  -------------
disk_name    system-1  (118065481)  Pool0  J8Y0TDZC       system-1
(118065481)
disk_name    system-1  (118065481)  Pool0  J8Y0TDZC       system-1
(118065481)
.
.
.
```

7. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

- Restore onboard key management encryption keys
- Restore external key management encryption keys

8. Boot the node: `boot_ontap`

Option 3: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `metrocluster node show -fields node-systemid,dr-partner-systemid`

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

```
dr-group-id cluster         node                 node-systemid dr-
partner-systemid
 ----------- --------------------- ------------------- -------------
-------------------
 1           Cluster_A             Node_A_1             536872914
118073209
 1           Cluster_B             Node_B_1             118073209
536872914
 2 entries were displayed.
```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

```
Local System ID: 118065481
    ...
    ...
```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                  POOL   SERIAL NUMBER  HOME
-------    -------------          -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Verify that the coredumps are saved: `system node run -node *local-node-name* partner savecore`

      If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node *local-node-name* partner savecore -s command`.</info>.

   c. Return to the admin privilege level: `set -privilege admin`

7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`

8. Boot the *replacement* node: `boot_ontap`

9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`

10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

       You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

    b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`
    c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - ASA A150

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.

4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the DIMM - ASA A150

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

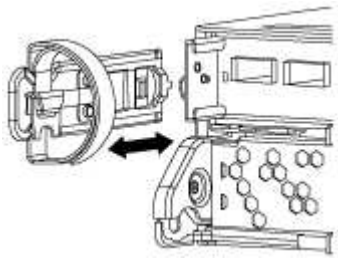

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

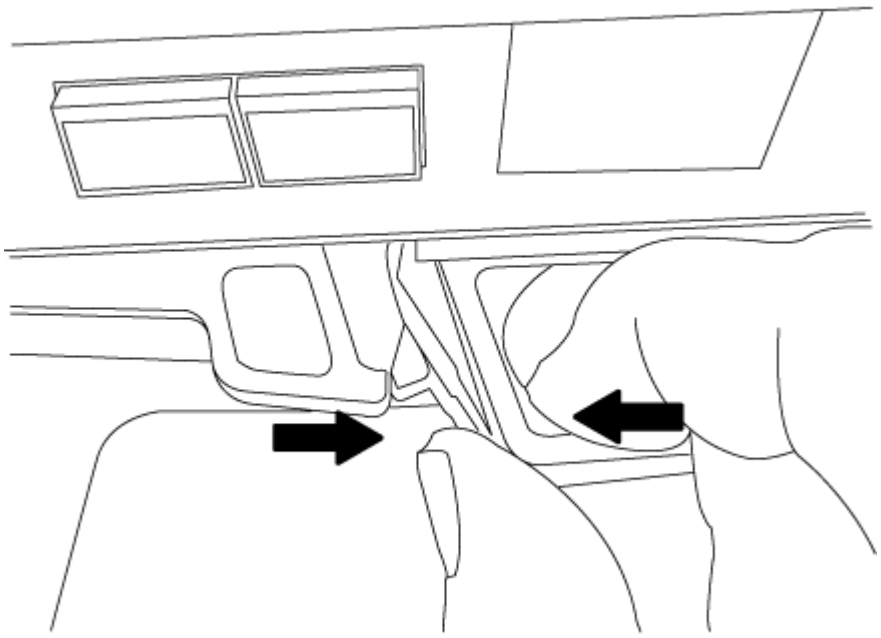

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

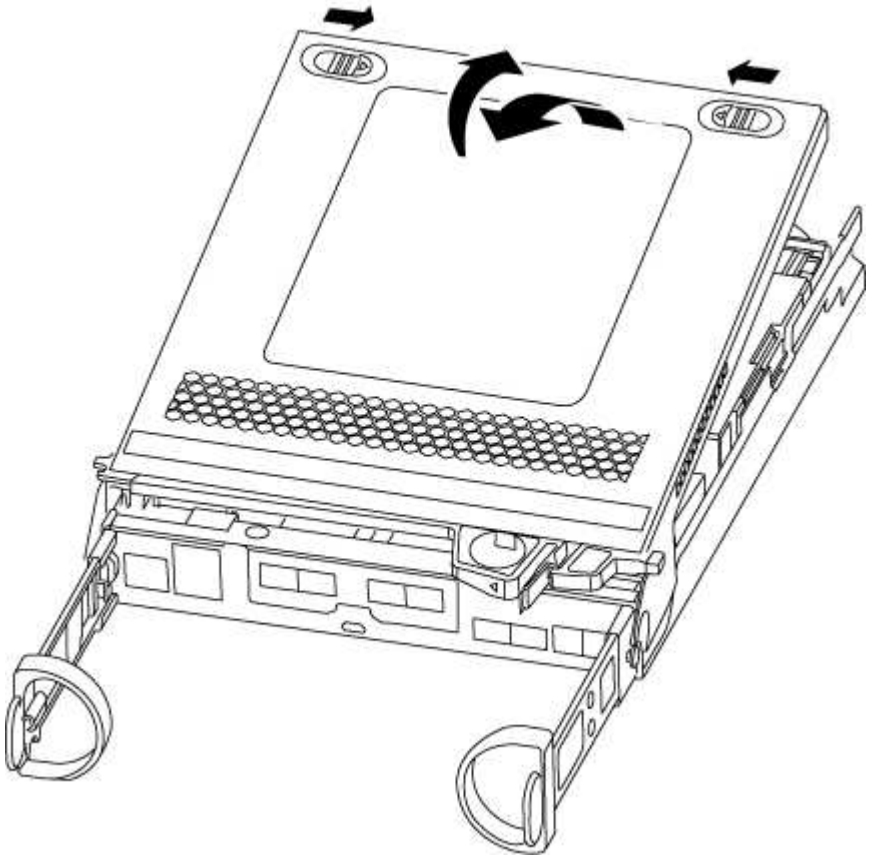

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

If you are replacing a DIMM, you need to remove it after you have unplugged the NVMEM battery from the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED on the back of controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

   

3. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
4. If the NVMEM LED is flashing, there is data in the NVMEM and you must disconnect the battery to clear the memory:
   a. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

   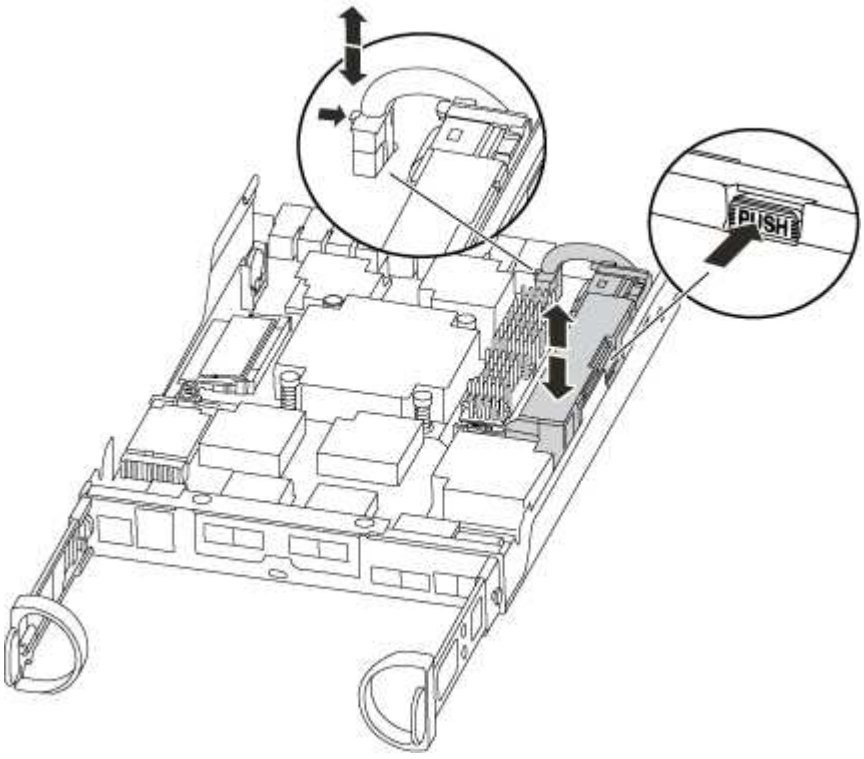

   b. Confirm that the NVMEM LED is no longer lit.
   c. Reconnect the battery connector.

5. Return to Step 3: Replace the DIMMs of this procedure to recheck the NVMEM LED.
6. Locate the DIMMs on your controller module.
7. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
8. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   The number and placement of system DIMMs depends on the model of your system.

   The following illustration shows the location of system DIMMs:

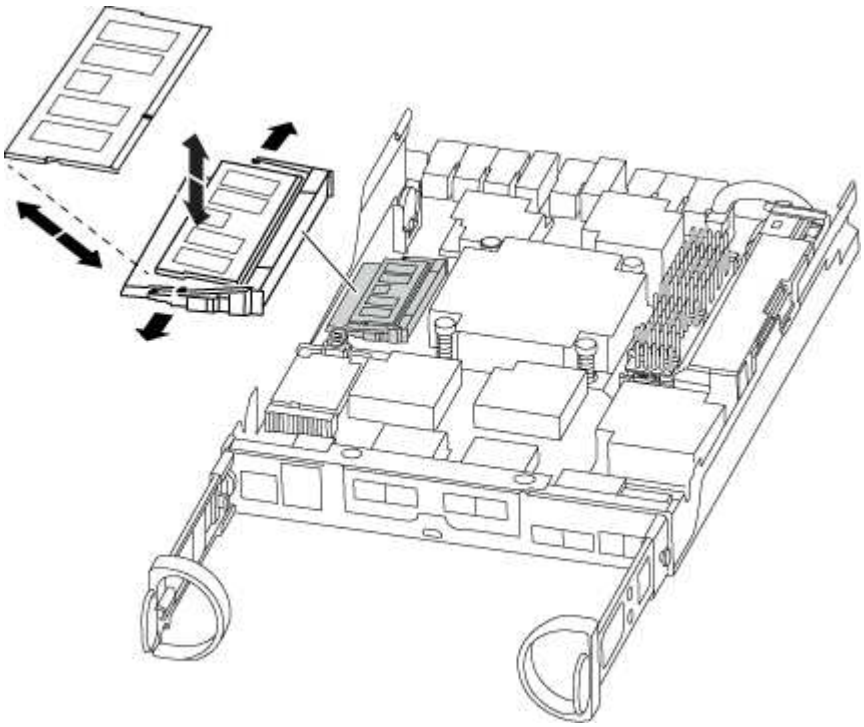

9. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

10. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

    The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

    Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

11. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
12. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

Make sure that the plug locks down onto the controller module.

13. Close the controller module cover.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, then turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the SSD drive or HDD Drive - ASA 150

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace the NVMEM battery - ASA A150

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

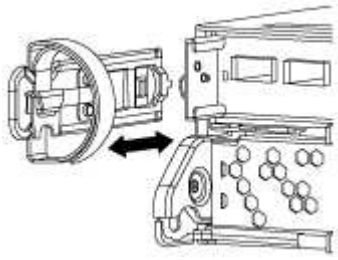

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

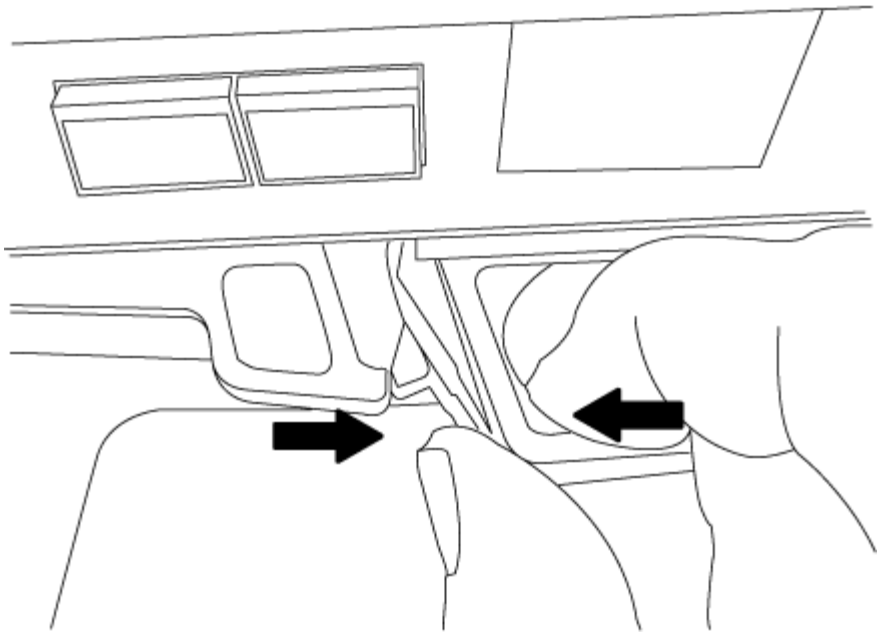

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery in your system, you must remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

3. Locate the NVMEM battery in the controller module.

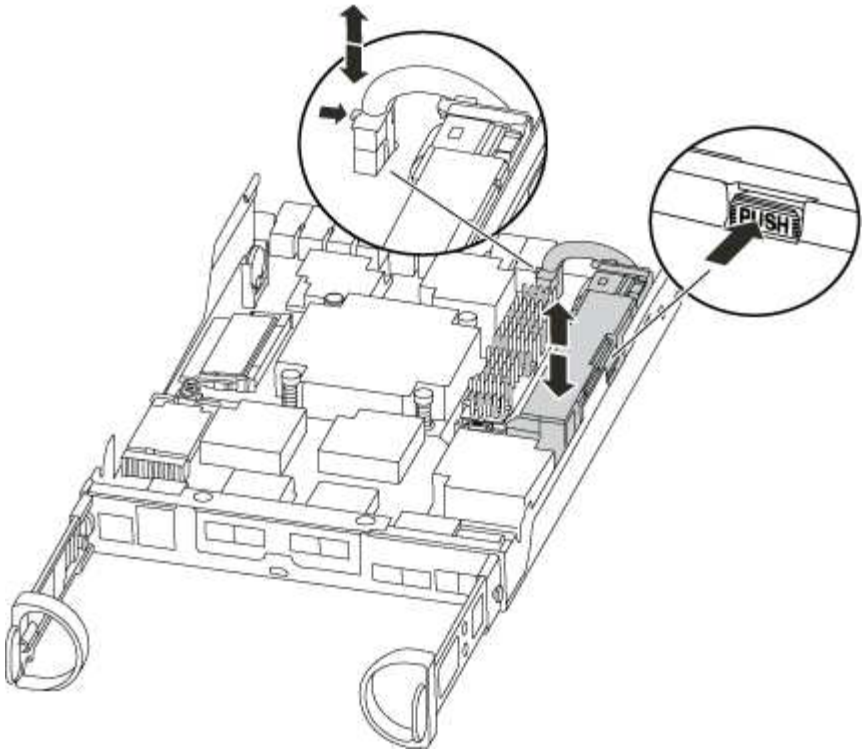

4. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
5. Remove the battery from the controller module and set it aside.
6. Remove the replacement battery from its package.
7. Loop the battery cable around the cable channel on the side of the battery holder.
8. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
9. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
10. Plug the battery plug back into the controller module.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber

optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - ASA A150

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> Cooling is integrated with the power supply, so you must replace the power supply within two minutes of removal to prevent overheating due to reduced airflow. Because the chassis provides a shared cooling configuration for the two HA nodes, a delay longer than two minutes will shut down all controller modules in the chassis. If both controller modules do shut down, make sure that both power supplies are inserted, turn both off for 30 seconds, and then turn both on.

- Power supplies are auto-ranging.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.

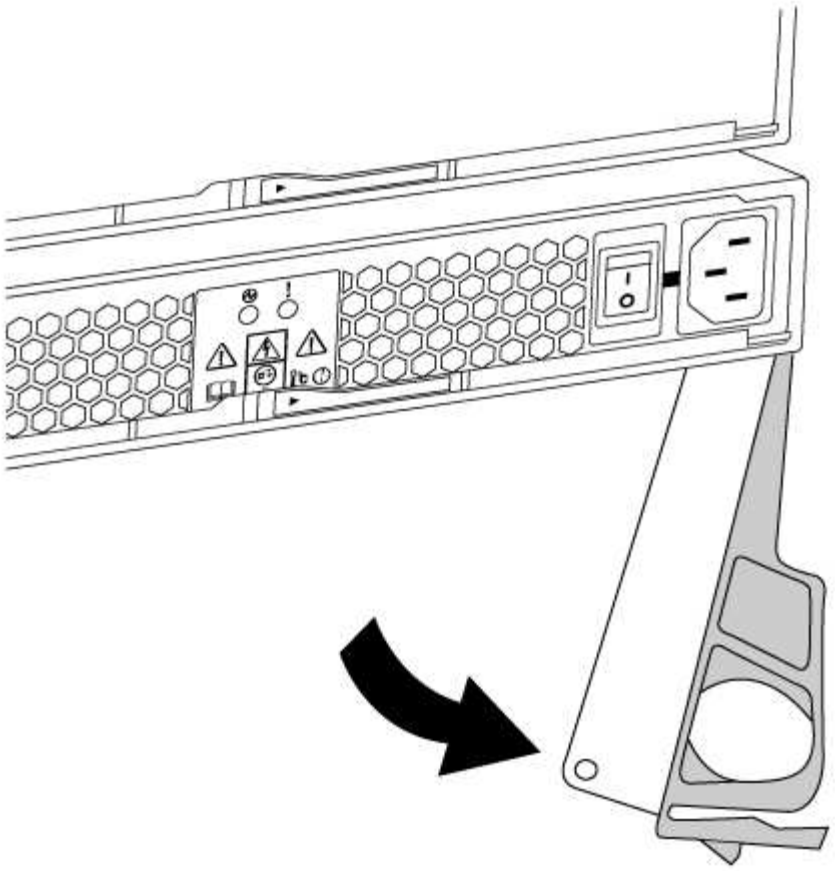

5. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.
7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
9. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.

10. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

    The power supply LEDs are lit when the power supply comes online.

11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA A150

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

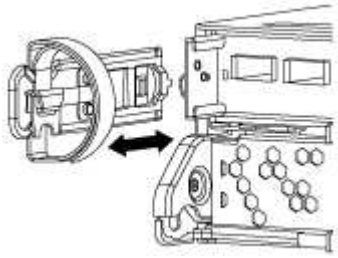

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

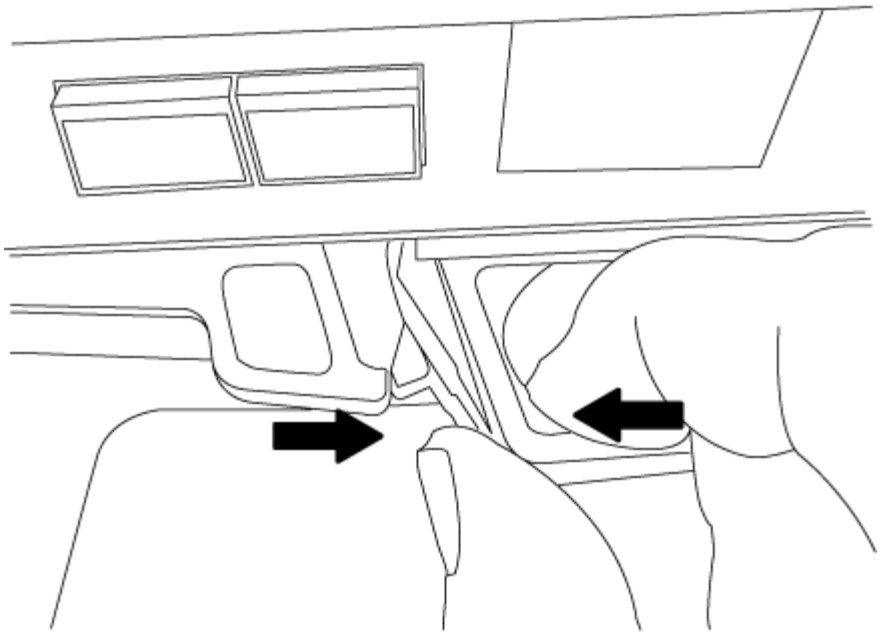

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

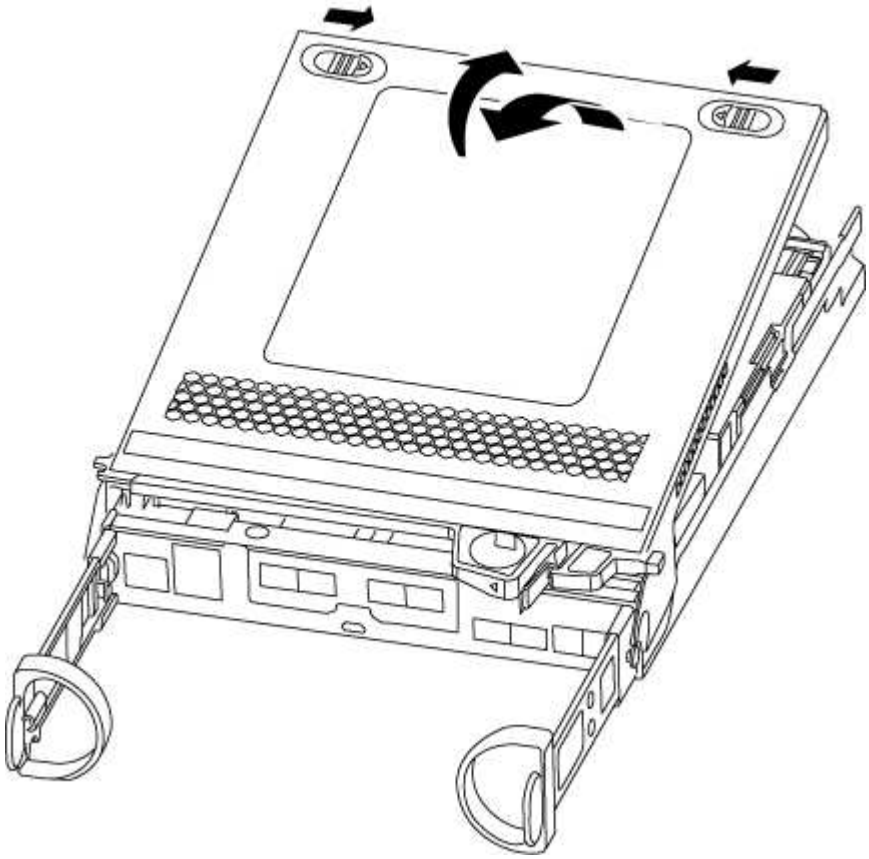

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

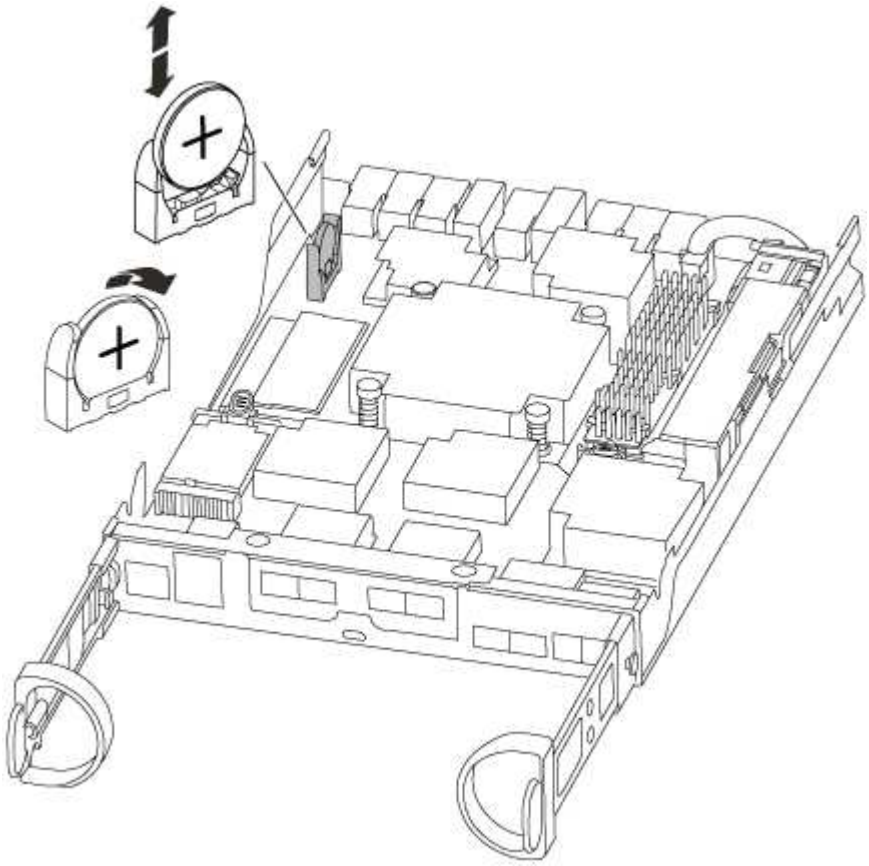

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and set time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA A250 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick steps - ASA A250

The Installation and Setup instructions give graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

The ASA A250 and ASA C250 use the same installation procedure as the AFF A250 system.

AFF A250 Installation and Setup Instructions

Video steps - ASA A250

The following video shows how to install and cable your new system.

Animation - Installation and Setup of an AFF A250

The ASA A250 uses the same installation procedure as the AFF A250 system.

Detailed steps - ASA A250

This page gives detailed step-by-step instructions for installing an ASA A250 system.

Step 1: Prepare for installation

To install your system, you need to create an account and register the system. You also need to inventory the

appropriate number and type of cables for your system and collect specific network information.

Customers with specific power requirements must check HWU for their configuration options.

Before you begin

- Make sure you have access to the NetApp Hardware Universe (HWU) for information about site requirements.
- Make sure you have access to the Release Notes for your version of ONTAP for more information about this system.
- You need to provide the following at your site:
  - Rack space for the storage system
  - Phillips #2 screwdriver
  - Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register your system.
4. Download and install NetApp Downloads: Config Advisor on your laptop.
5. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
25 GbE cable	X66240A-05 (112-00595), 0.5m; X66240-2 (112-00573), 2m		Cluster interconnect network
	X66240A-2 (112-00598), 2m; X66240A-5 (112-00600), 5m		Data
100 GbE cable	X66211-2 (112-00574), 2m; X66211-5 (112-00576), 5m		Storage

Type of cable…	Part number and length	Connector type	For…
RJ-45 (order dependent)	Not applicable		Management network (BMC and wrench port) and Ethernet data (e0a and e0b)
Fibre Channel	X66250-2 (112-00342) 2m; X66250-5 (112-00344) 5m; X66250-15 (112-00346) 15m; X66250-30 (112-00347) 30m		
Micro-USB console cable	Not applicable		Console connection during software setup
Power cables	Not applicable		Powering up the system

6. Review the ONTAP Configuration Guide and collect the required information listed in that guide.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

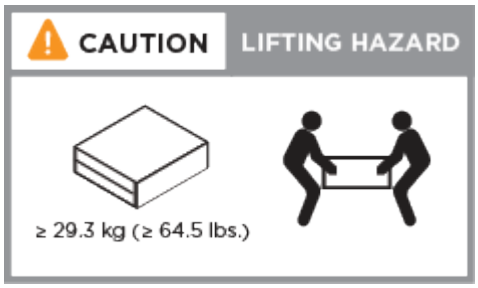

3. Identify and manage cables because this system does not have a cable management device.
4. Place the bezel on the front of the system.

Step 3: Cable controllers to cluster

Cable the controllers to a cluster by using the two-node switchless cluster method or by using the cluster interconnect network method.

Option 1: Two-node switchless cluster

The management, Fibre Channel, and data or host network ports on the controller modules are connected to switches. The cluster interconnect ports are cabled on both controller modules.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

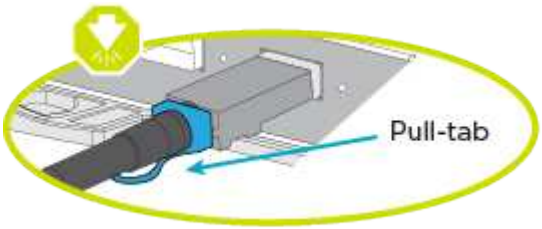

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the tabulated steps to complete the cabling between the controllers and the switches. Perform the steps on each controller.

Animation - Cable two-node switchless cluster

Steps

1. Use the the 25GbE cluster interconnect cable to connect the cluster interconnect ports e0c to e0c and e0d to e0d.


:

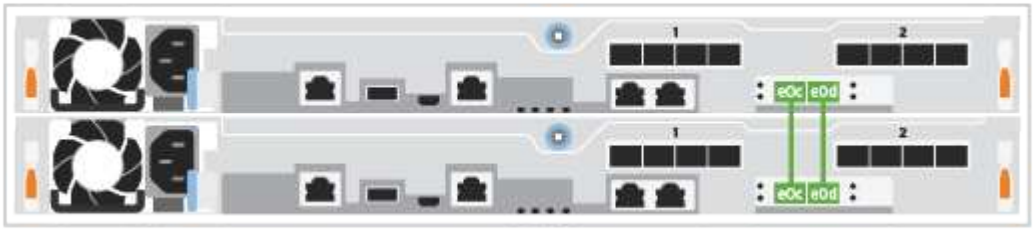

2. Cable the wrench ports to the management network switches with the RJ45 cables.

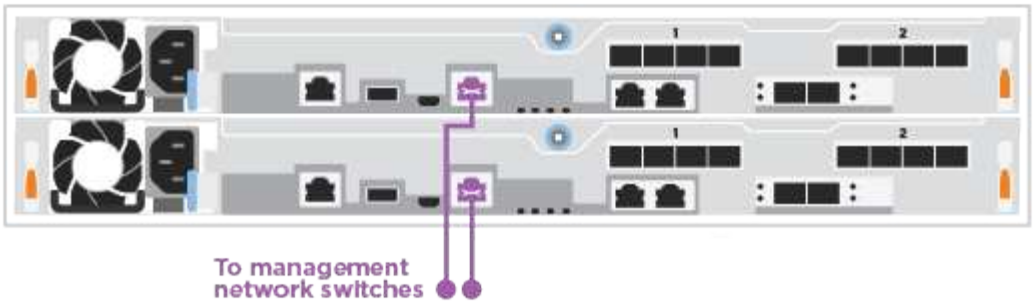

DO NOT plug in the power cords at this point.

Option 2: Switched cluster

All ports on the controllers are connected to switches; cluster interconnect, management, Fibre Channel, and data or host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

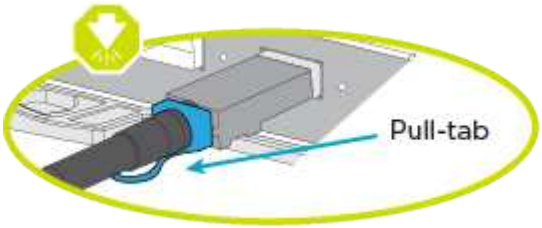

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the steps to complete the cabling between the controllers and the switches. Perform the steps on each controller.

Animation - Cable switched cluster

Steps

1. Cable the cluster interconnect ports e0c and e0d to the 25 GbE cluster interconnect switches.

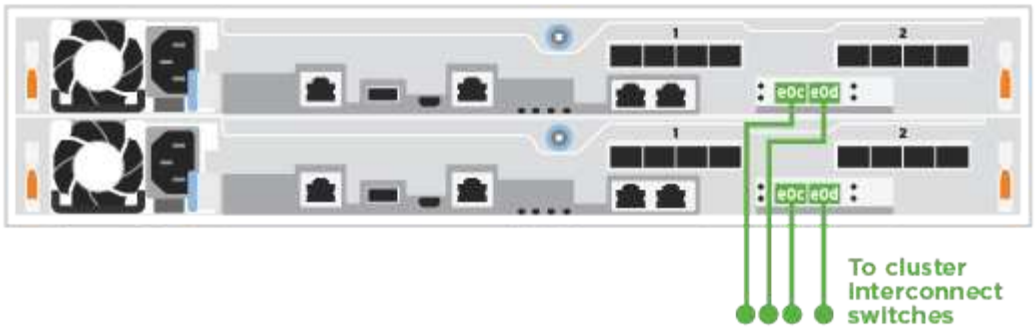

2. Cable the wrench ports to the management network switches with the RJ45 cables.

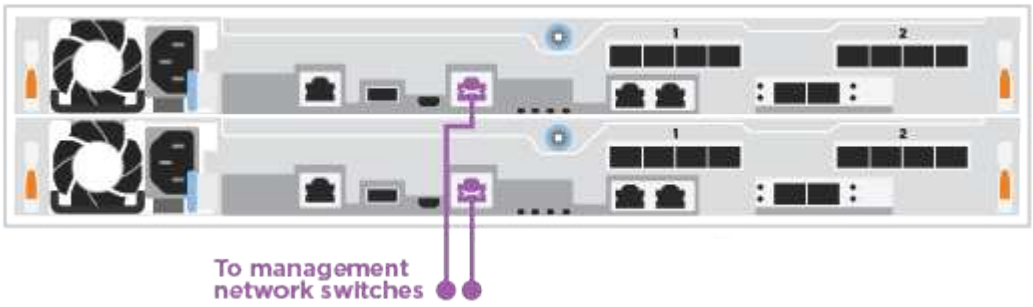

Step 4: Cable to host network or storage (Optional)

You have configuration-dependent optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cabling to a host network and storage.

Option 1: Cable to Fibre Channel host network

Fibre Channel ports on the controllers are connected to Fibre Channel host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

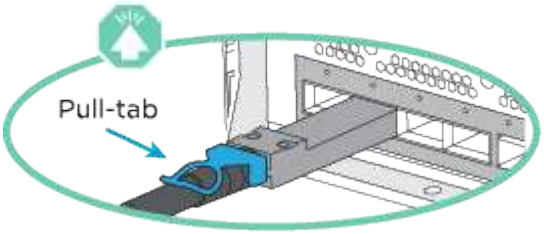

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again..

About this task

Perform the following step on each controller module.

Steps

1. Cable ports 2a through 2d to the FC host switches.

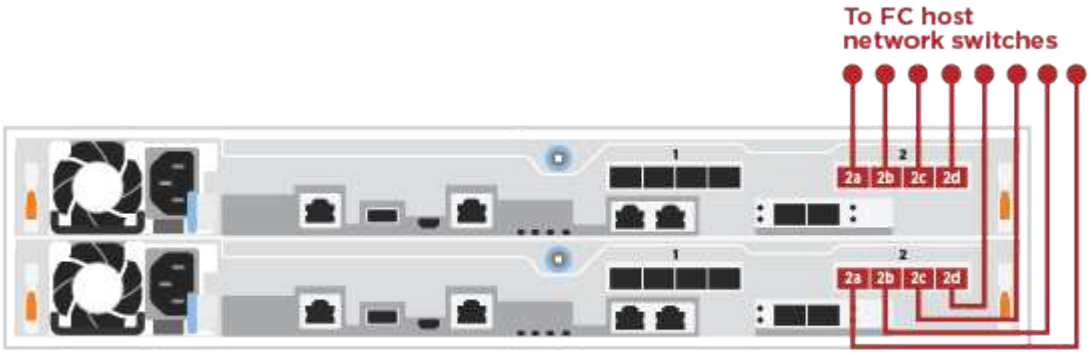

Option 2: Cable to 25GbE data or host network

25GbE ports on the controllers are connected to 25GbE data or host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

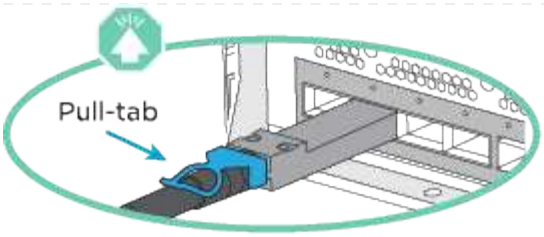

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Perform the following step on each controller module.

Steps

1. Cable ports e4a through e4d to the 10GbE host network switches.

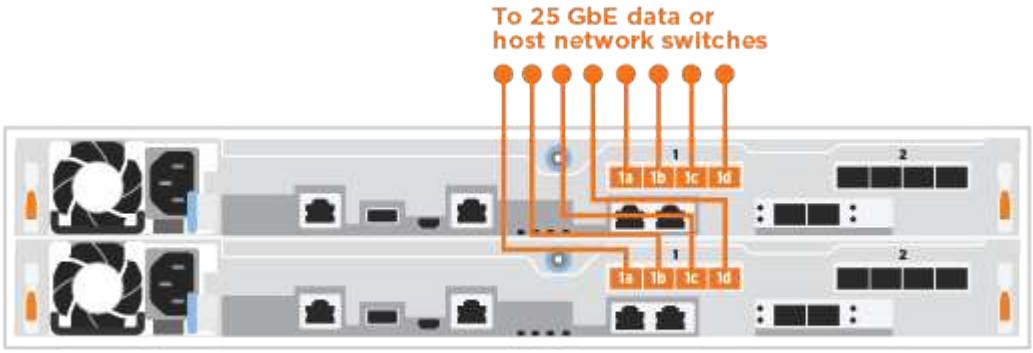

Option 3: Cable controllers to single drive shelf

Cable each controller to the NSM modules on the NS224 drive shelf.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

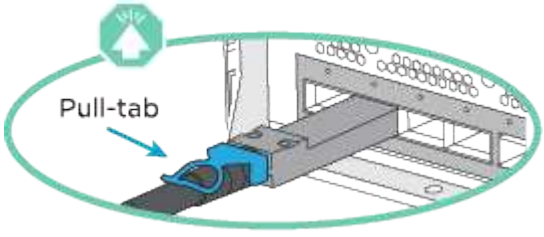

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the tabulated steps to complete the cabling between the controllers and the single shelf. Perform the steps on each controller module.

Animation - Cable the controllers to a single NS224

Steps

1. Cable controller A to the shelf.

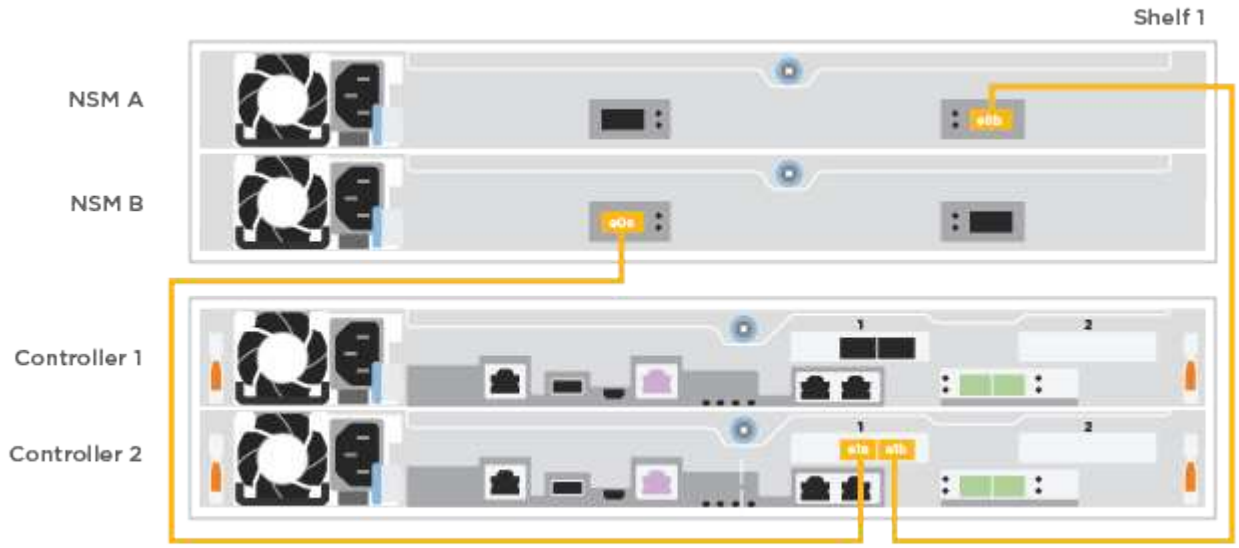

2. Cable controller B to the shelf.

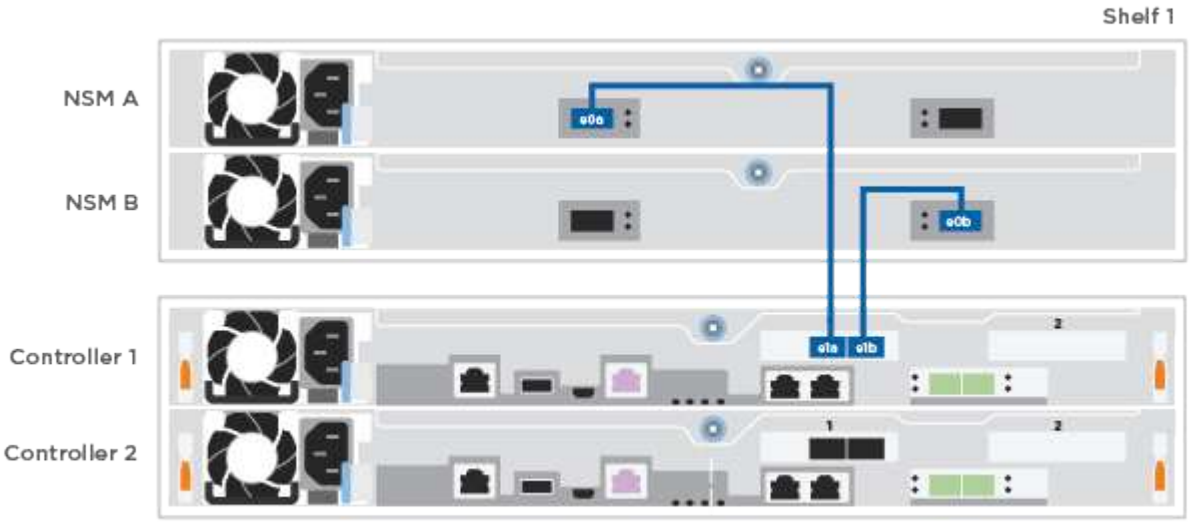

Step 5: Complete system setup

Complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: If network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

2. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

3. Use the animation to connect your laptop to the Management switch:

   Animation - Connect your laptop to the Management switch

4. Select an ONTAP icon listed to discover:

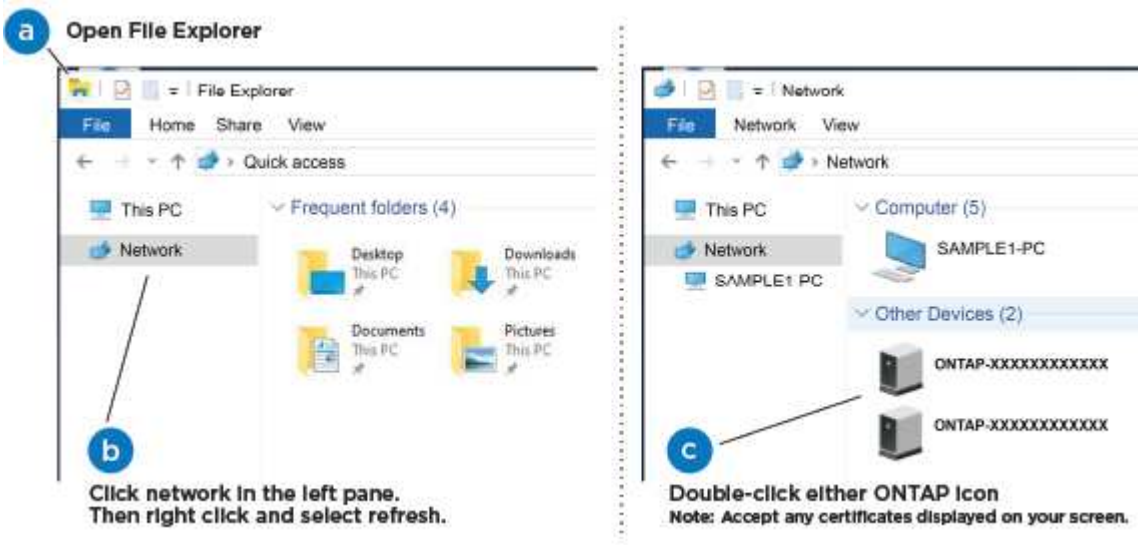

   a. Open File Explorer.

   b. Click **Network** in the left pane and right-click and select **refresh**.

   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

   XXXXX is the system serial number for the target node.

   System Manager opens.

5. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
6. Verify the health of your system by running Config Advisor.
7. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the laptop or console to the switch on the management subnet.

   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.
3. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

4. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the ONTAP Configuration Guide.

5. Verify the health of your system by running Config Advisor.
6. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain ASA A250 hardware

For the ASA A250 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

Fan

The fan cools the controller.

Mezzanine card

A Mezzanine card is a printed circuit board that plugs directly into another plug-in card.

NVEM battery

A battery is included with the controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview and requirements ASA A250

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.
- You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

About this task

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
  - The *impaired* node is the controller on which you are performing maintenance.
  - The *healthy* node is the HA partner of the impaired controller.

Check encryption key support and status - ASA A250

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select

one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller - ASA A250

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller displays…	Then…
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Option 2: Systems in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *impaired_node_name* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - ASA A250

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller module, you must first remove the controller module from the system, and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Unplug the I/O cables from the controller module.
5. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

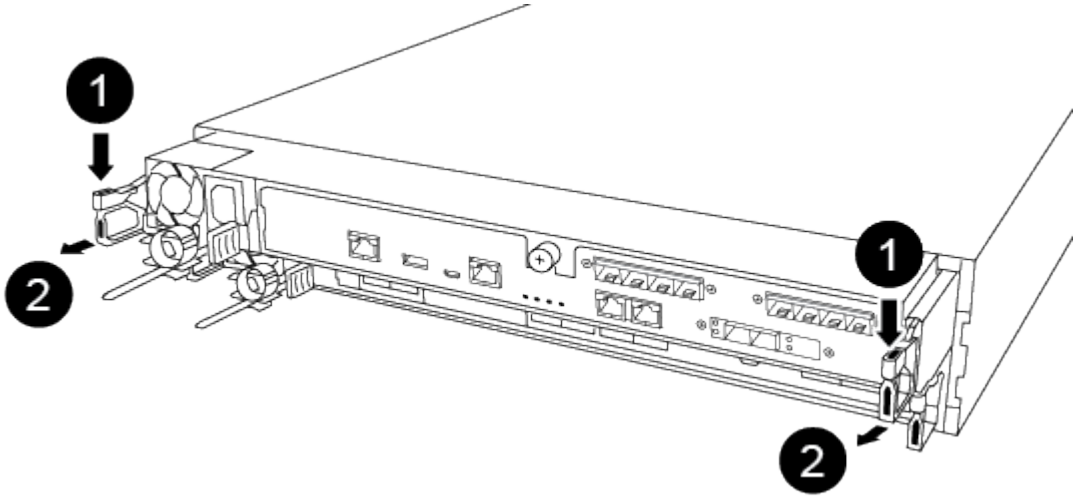

1	Lever
2	Latching mechanism

6. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
7. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

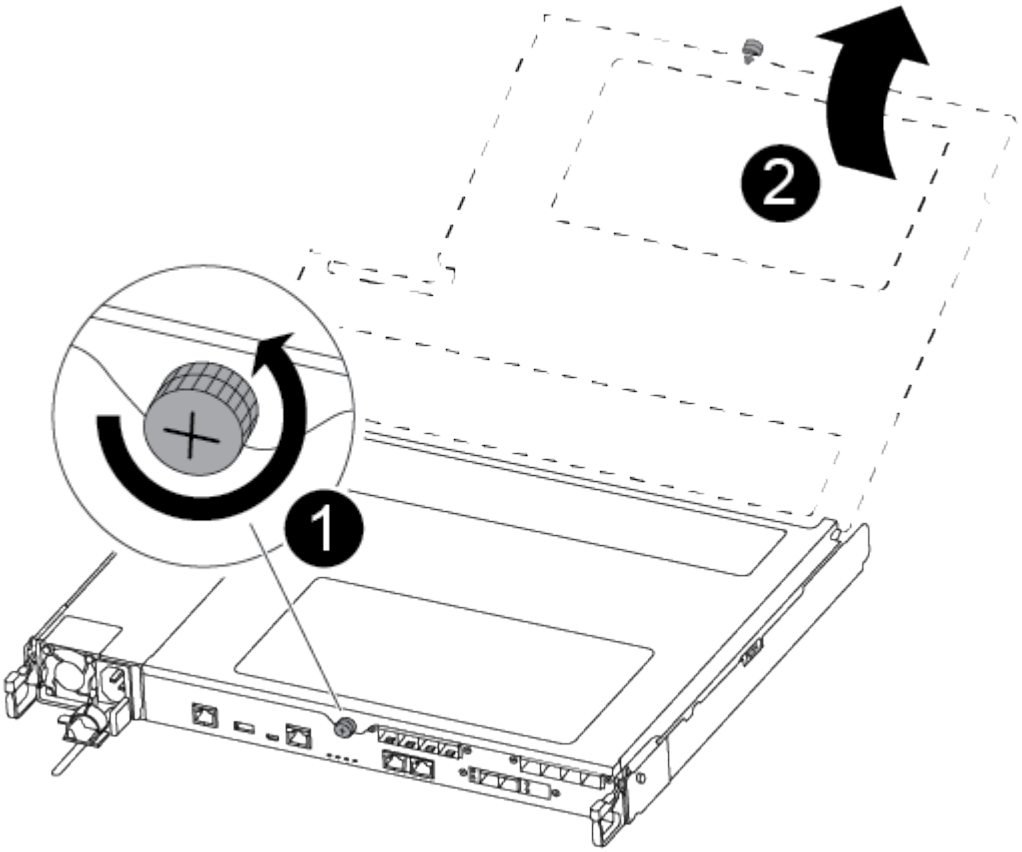

1	Thumbscrew
2	Controller module cover.

8. Lift out the air duct cover.

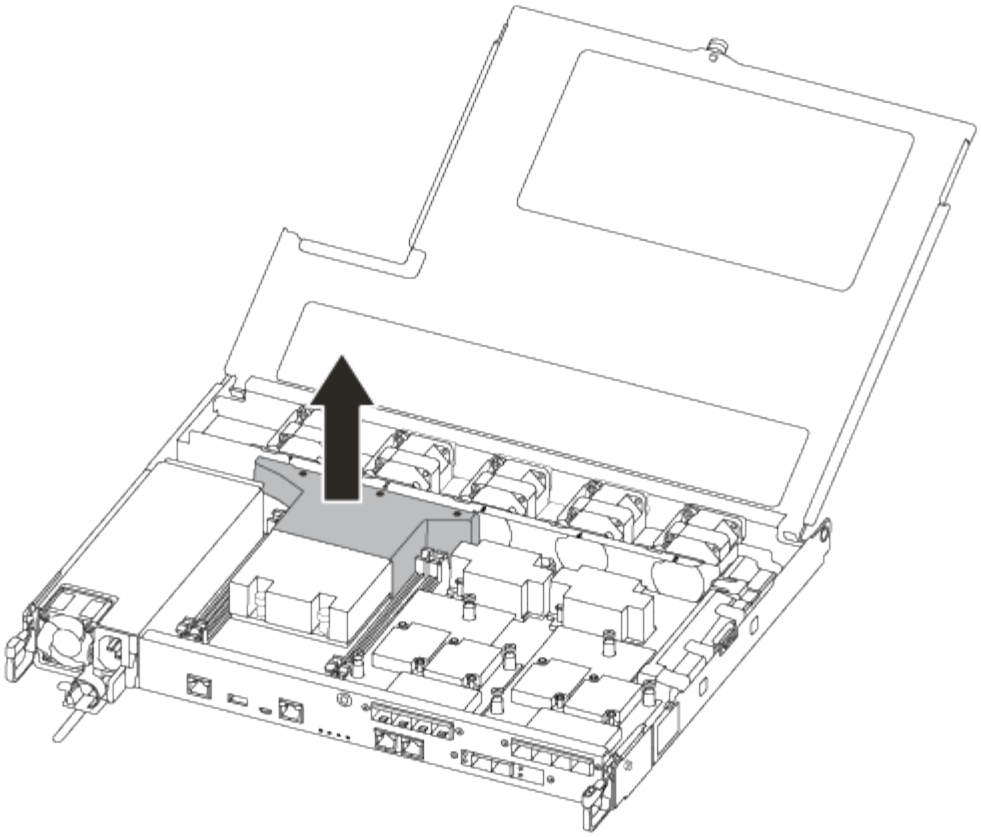

Step 2: Replace the boot media

You locate the failed boot media in the controller module by removing the air duct on the controller module before you can replace the boot media.

You need a #1 magnetic Phillips head screwdriver to remove the screw that holds the boot media in place. Due to the space constraints within the controller module, you should also have a magnet to transfer the screw on to so that you do not lose it.

You can use the following video or the tabulated steps to replace the boot media:

Animation - Replace the boot media

1. Locate and replace the impaired boot media from the controller module.

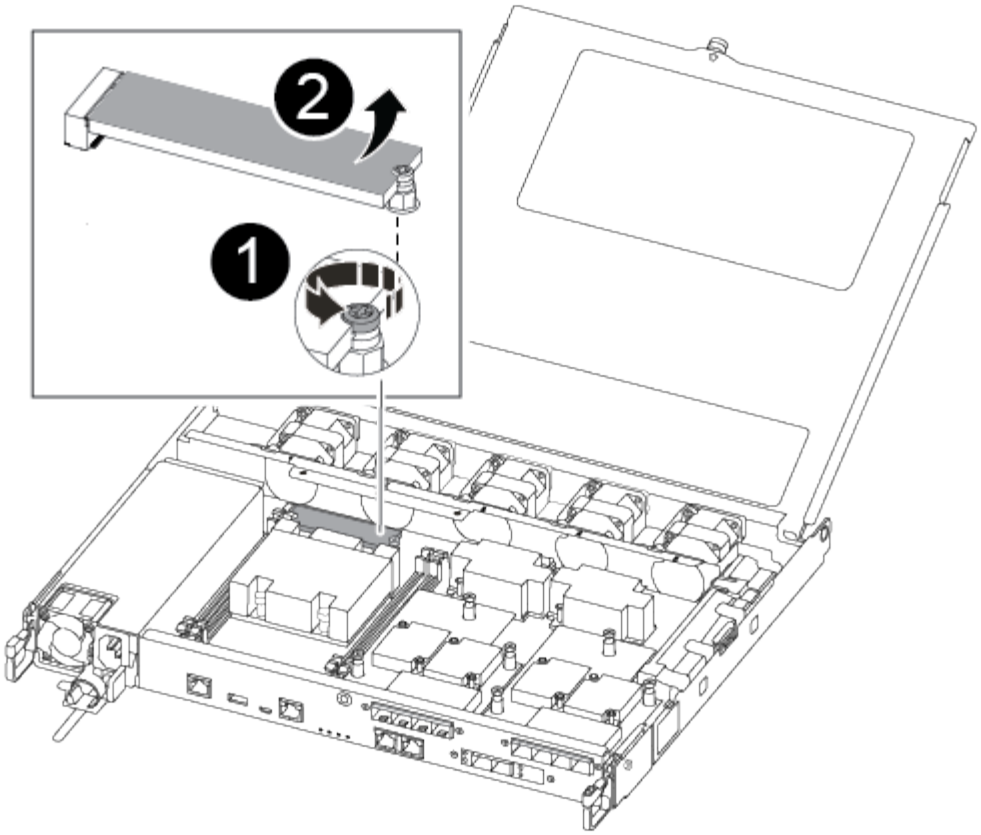

1	Remove the screw securing the boot media to the motherboard in the controller module.
2	Lift the boot media out of the controller module.

2. Using the #1 magnetic screwdriver, remove the screw from the impaired boot media, and set it aside safely on the magnet.
3. Gently lift the impaired boot media directly out of the socket and set it aside.
4. Remove the replacement boot media from the antistatic shipping bag and align it into place on the controller module.
5. Using the #1 magnetic screwdriver, insert and tighten the screw on the boot media.

> Do not apply force when tightening the screw on the boot media; you might crack it.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed is without a boot image so you need to transfer a boot image using a USB flash drive.

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download

button.

- If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
2. Download the service image to your work space on your laptop.
3. Unzip the service image.

   If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

   There are two folders in the unzipped service image file:

   - boot
   - efi

4. Copy the efi folder to the top directory on the USB flash drive.

   If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

   The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

5. Remove the USB flash drive from your laptop.
6. If you have not already done so, install the air duct.

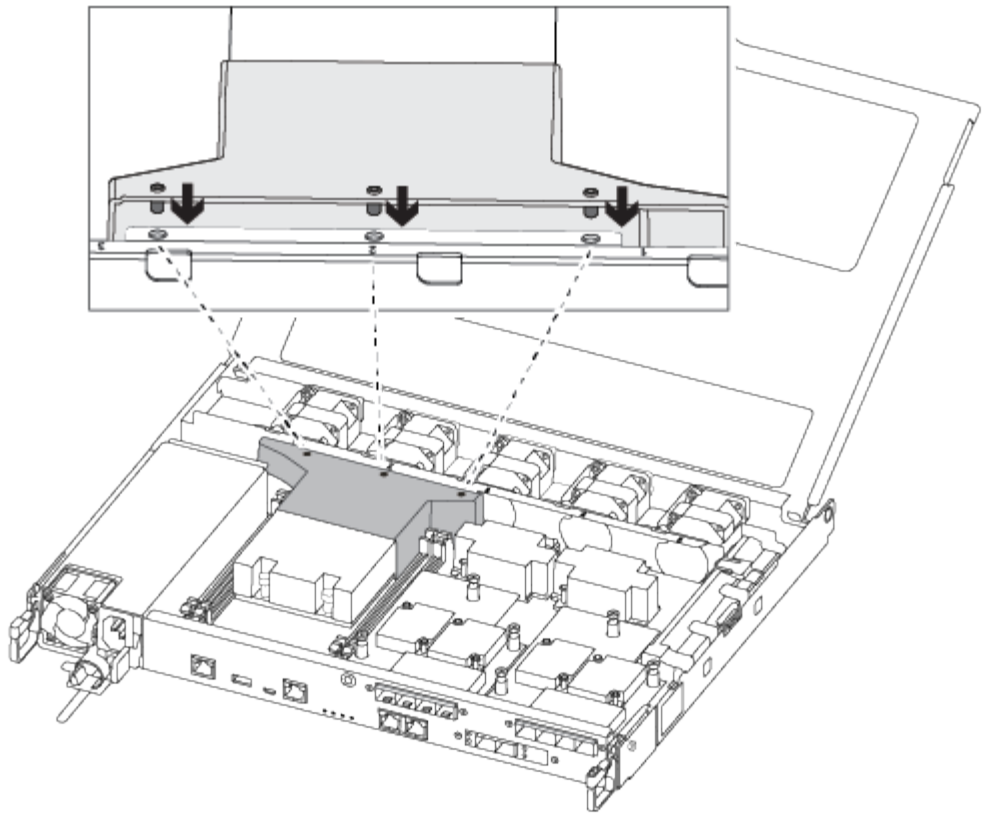

7. Close the controller module cover and tighten the thumbscrew.

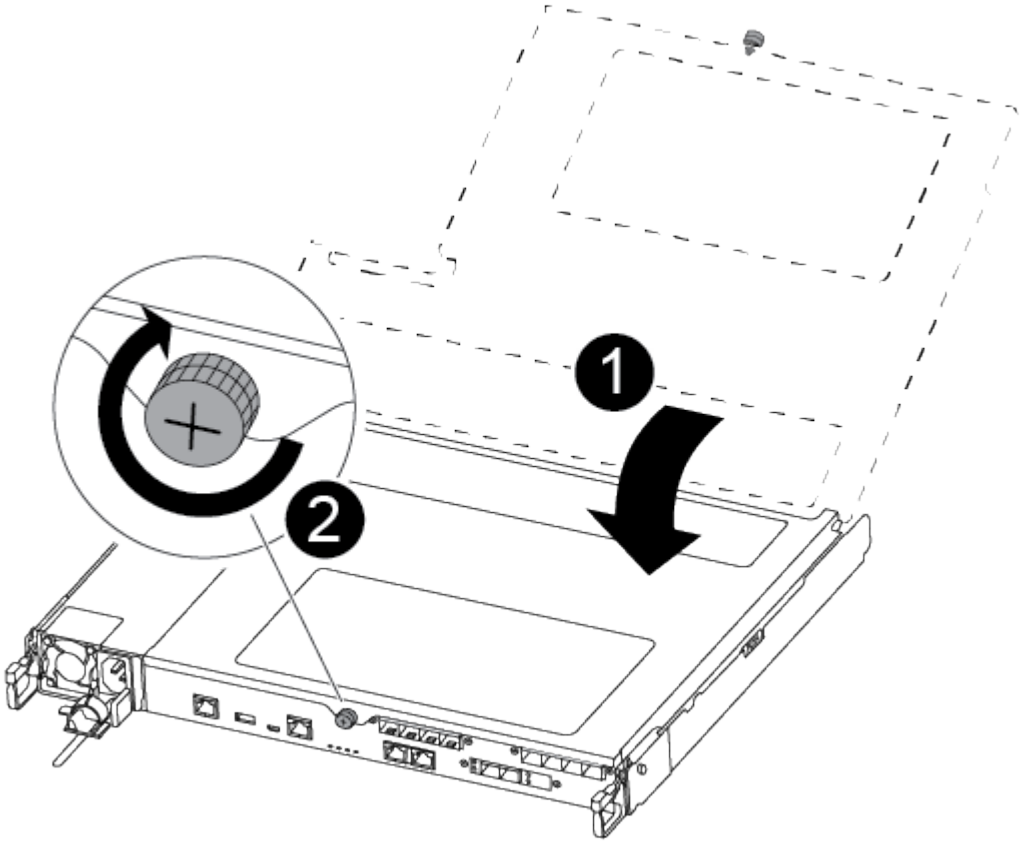

1	Controller module cover
2	Thumbscrew

8. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
9. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

10. Push the controller module all the way into the chassis:
11. Place your index fingers through the finger holes from the inside of the latching mechanism.
12. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
13. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

    The controller module should be fully inserted and flush with the edges of the chassis.

14. Reconnect the controller module I/O cables.
15. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

    The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

16. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

    If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

17. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

    The system begins to boot and stops at the LOADER prompt.

Boot the recovery image - ASA A250

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.

3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.

5. Give back the controller using the `storage failover giveback -fromnode local` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

   NOTE: If the process fails, contact NetApp Support.

Restore encryption - ASA A250

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
**********************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
**********************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
**********************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
**********************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - ASA A250

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - ASA A250

To replace the chassis, you must move the bezel, controller modules, and NVMe drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in its place.

About this task

- All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - ASA A250

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

```
Warning: Are you sure you want to halt node <node_name>? {y|n}:
```

10. Wait for each controller to halt and display the LOADER prompt.

Replace chassis - ASA A250

To replace the chassis, you move the power supplies, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis from with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

Use the following video or the tabulated steps to replace the chassis; it assumes the removal and replacement of the bezel:

Animation - Replace the chassis

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

> If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

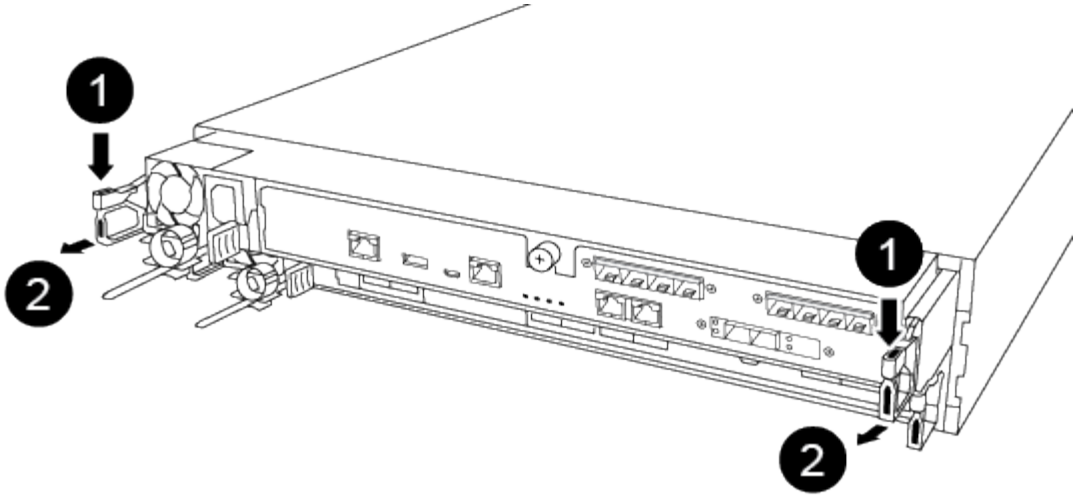

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.
5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up

and against the drive holder.

Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot the system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Insert the controller module into the chassis:

   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

   The controller module should be fully inserted and flush with the edges of the chassis.

   a. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect

the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - ASA A250

You must verify the HA state of the chassis, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis *HA-state*`

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement- ASA A250

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct

recovery procedure to determine whether you should use this procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller module - ASA A250

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace controller - ASA A250

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

Use the following video or the tabulated steps to replace a controller module:

Animation - Replace a controller module

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

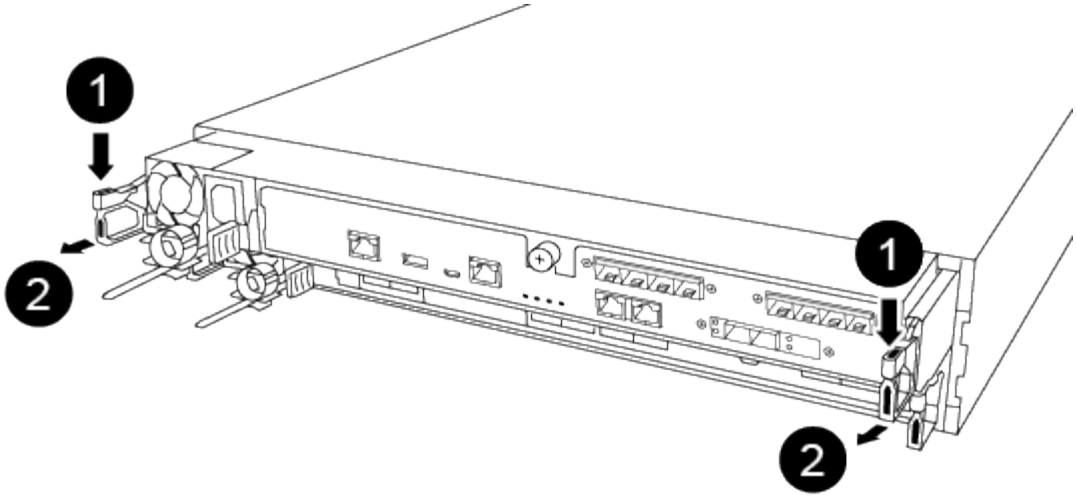

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

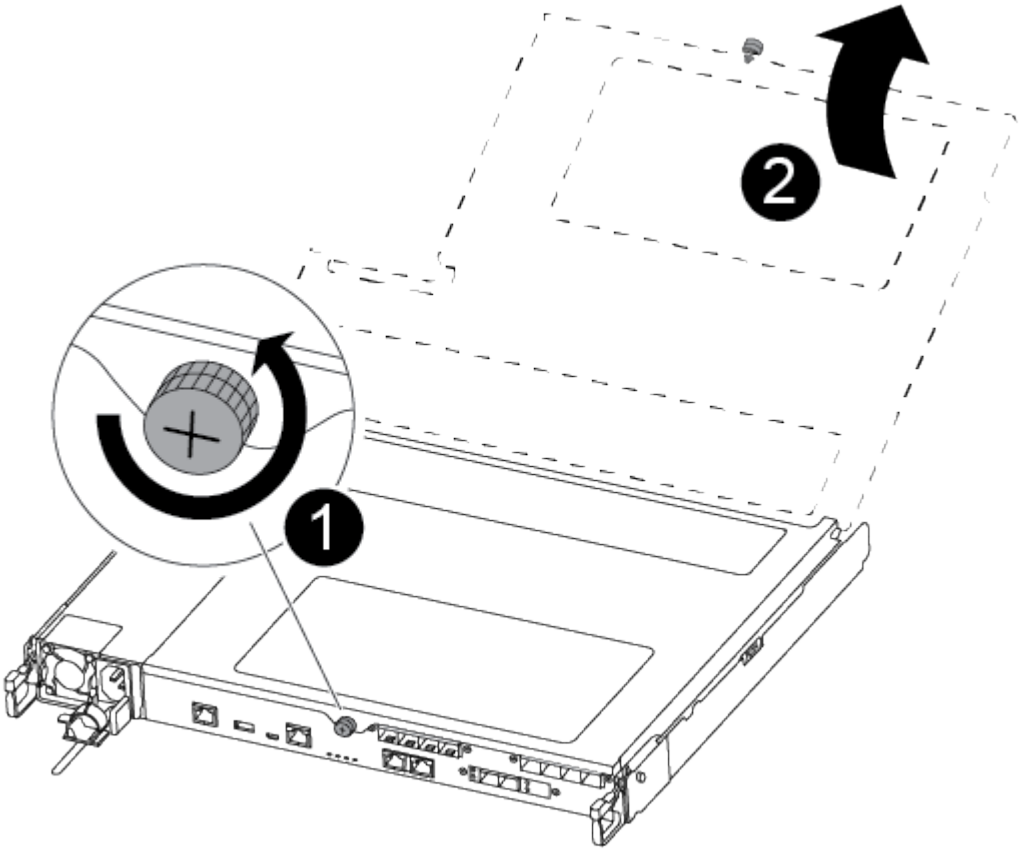

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

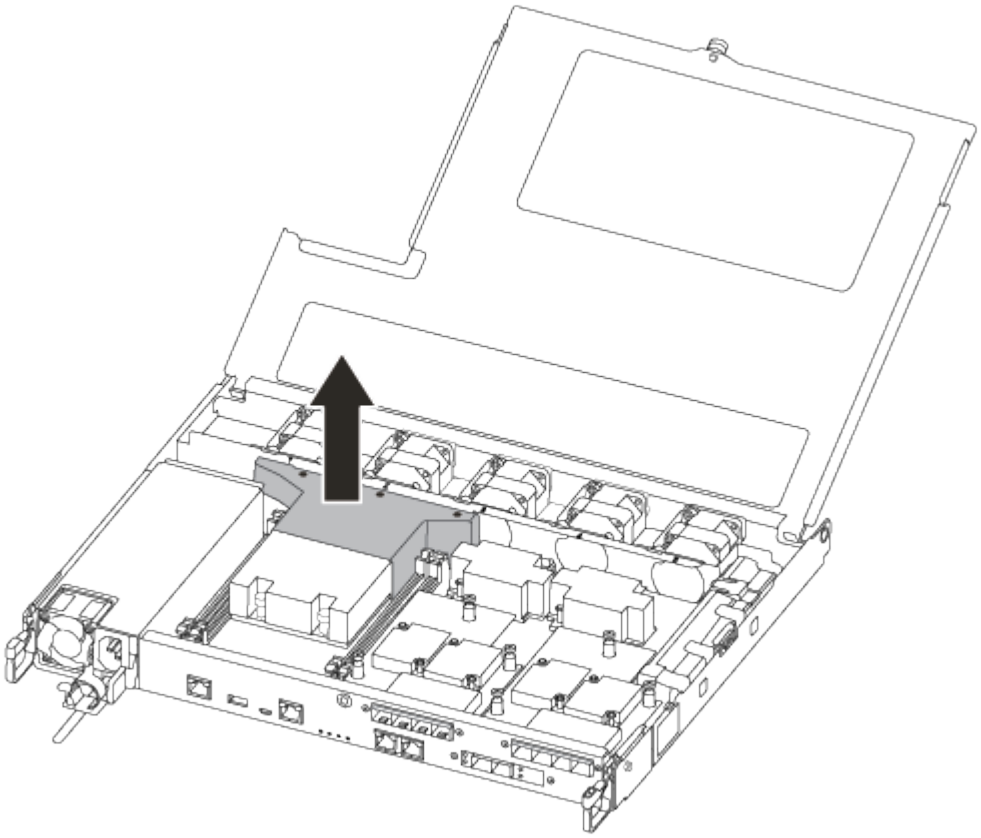

Step 2: Move the power supply

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

1. Disconnect the power supply.
2. Open the power cable retainer, and then unplug the power cable from the power supply.
3. Unplug the power cable from the power source.
4. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

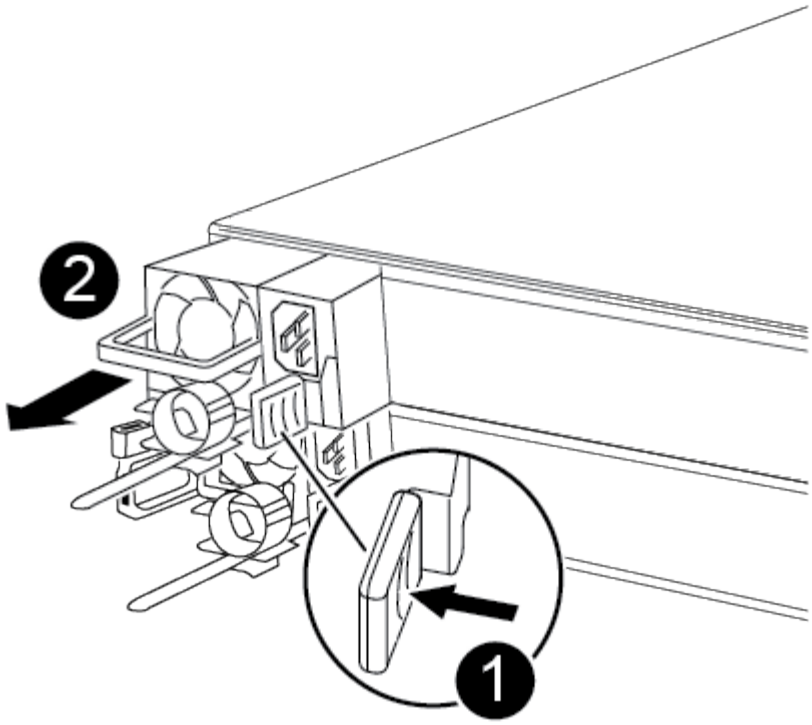

1	Blue power supply locking tab
2	Power supply

5. Move the power supply to the new controller module, and then install it.
6. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

You must move the fans from the impaired controller module to the replacement module when replacing a failed controller module.

1. Remove the fan module by pinching the side of the fan module, and then lifting the fan module straight out of the controller module.

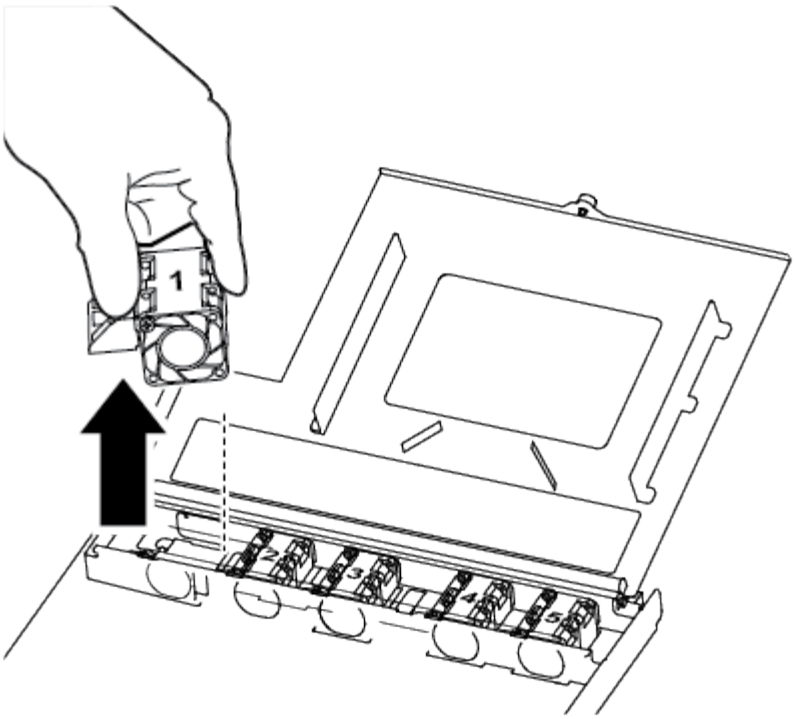

1	Fan module

2. Move the fan module to the replacement controller module, and align the edges of the fan module with the opening in the controller module, and then slide the fan module in.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the boot media

You must move the boot media device from the impaired controller module to the replacement controller module.

You need a #1 magnetic Phillips head screwdriver to remove the screw that holds the boot media in place. Due to the space constraints within the controller module, you should also have a magnet to transfer the screw on to so that you do not lose it.

1. Locate and move the boot media from the impaired controller module to the replacement controller module.

   The boot media is located under the air duct cover you removed earlier in this procedure.

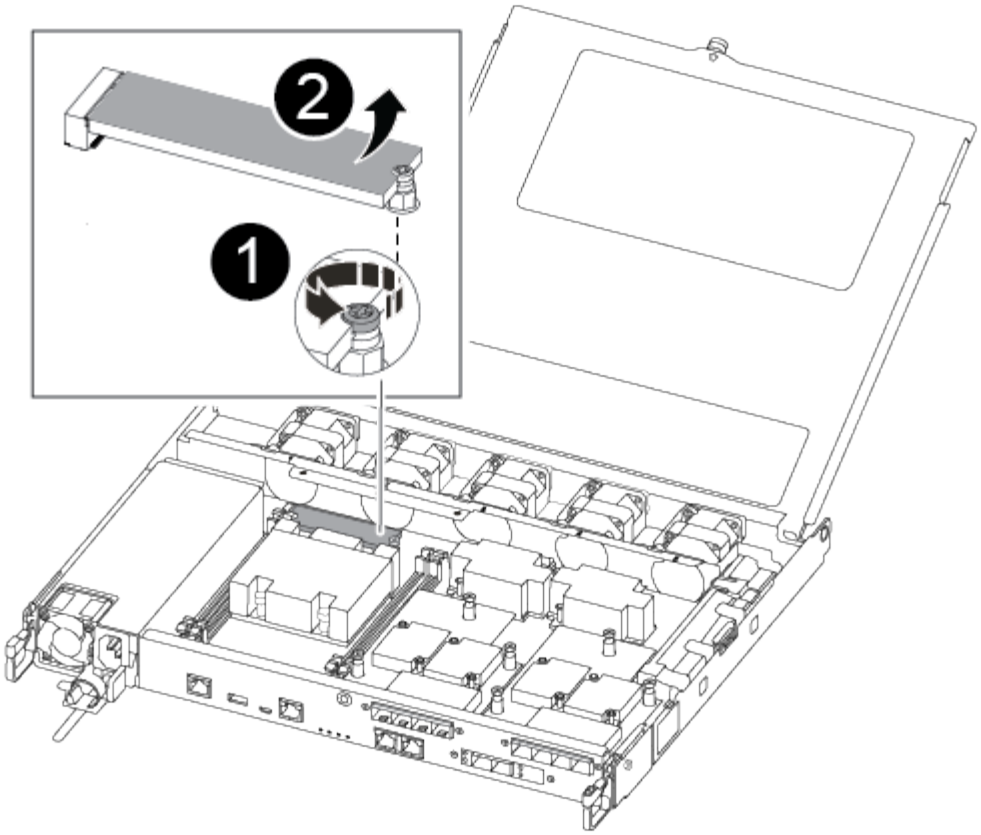

1	Remove the screw securing the boot media to the motherboard in the impaired controller module.
2	Lift the boot media out of the impaired controller module.

2. Using the #1 magnetic screwdriver, remove the screw from the boot media, and set it aside safely on the magnet.
3. Gently lift the boot media directly out of the socket and align it into place in the replacement controller module.
4. Using the #1 magnetic screwdriver, insert and tighten the screw on the boot media.

> Do not apply force when tightening the screw on the boot media; you might crack it.

Step 5: Move the DIMMs

To move the DIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

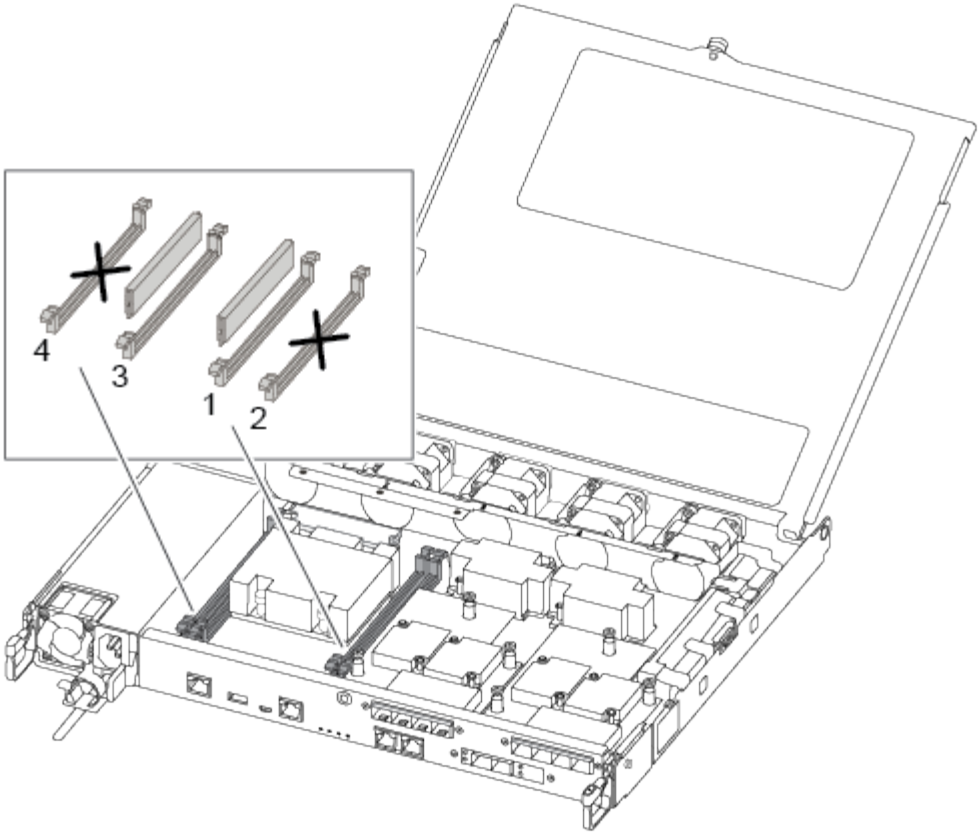

Install each DIMM into the same slot it occupied in the impaired controller module.

1. Slowly push apart the DIMM ejector tabs on either side of the DIMM, and slide the DIMM out of the slot.

   Hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

2. Locate the corresponding DIMM slot on the replacement controller module.
3. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the DIMM squarely into the socket.

   The DIMMs fit tightly in the socket. If not, reinsert the DIMM to realign it with the socket.

4. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.
5. Repeat these steps for the remaining DIMM.

Step 6: Move a mezzanine card

To move a mezzanine card, you must remove the cabling and any QSFPs and SFPs from the ports, move the mezzanine card to the replacement controller, reinstall any QSFPs and SFPs onto the ports, and cable the ports.

1. Locate and move the mezzanine cards from your impaired controller module.

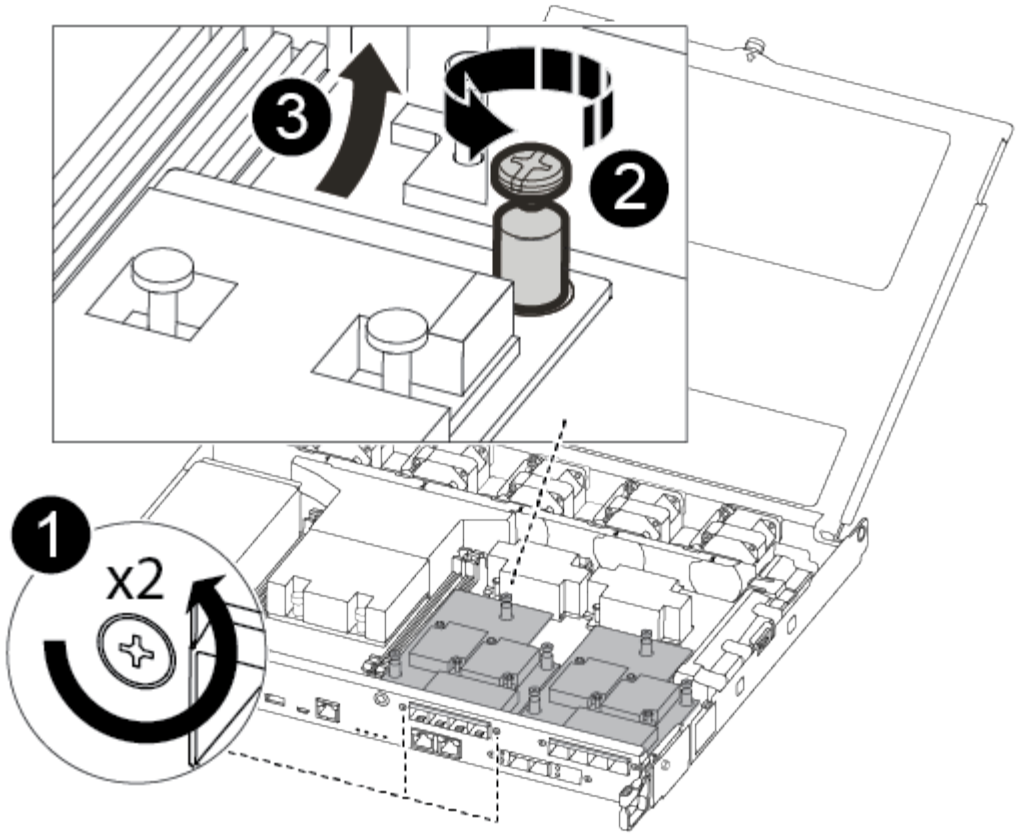

1	Remove screws on the face of the controller module.
2	Loosen the screw in the controller module.
3	Move the mezzanine card.

2. Unplug any cabling associated with the mezzanine card.

   Make sure that you label the cables so that you know where they came from.

   a. Remove any SFP or QSFP modules that might be in the mezzanine card and set it aside.
   b. Using the #1 magnetic screwdriver, remove the screws from the face of the impaired controller module and from the mezzanine card, and set them aside safely on the magnet.
   c. Gently lift the mezzanine card out of the socket and move it to the same position in the replacement controller.
   d. Gently align the mezzanine card into place in the replacement controller.
   e. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the replacement controller module and on the mezzanine card.

      Do not apply force when tightening the screw on the mezzanine card; you might crack it.

3. Repeat these steps if there is another mezzanine card in the impaired controller module.

4. Insert the SFP or QSFP modules that were removed onto the mezzanine card.

Step 7: Move the NV battery

When replacing the controller module, you must move the NV battery from the impaired controller module to the replacement controller module.

1. Locate and move the NVMEM battery from your impaired controller module to the replacement controller module.

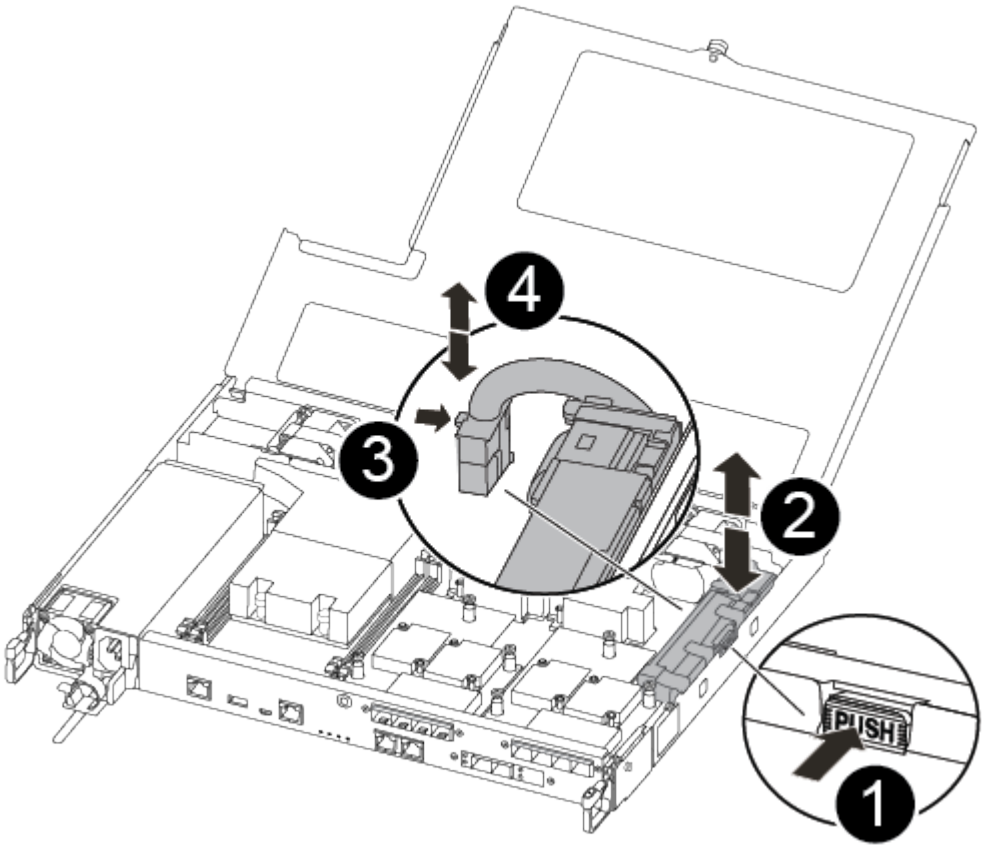

1	Squeeze the clip on the face of the battery plug.
2	Unplug the battery cable from the socket.
3	Grasp the battery and press the blue locking tab marked PUSH.
4	Lift the battery out of the holder and controller module.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.

4. Locate the corresponding NV battery holder on the replacement controller module and align the NV battery to the battery holder.
5. Insert the NV battery plug into the socket.
6. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
7. Press firmly down on the battery pack to make sure that it is locked into place.

Step 8: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

You can use the following illustrations or the written steps to install the replacement controller module in the chassis.

1. If you have not already done so, install the air duct.

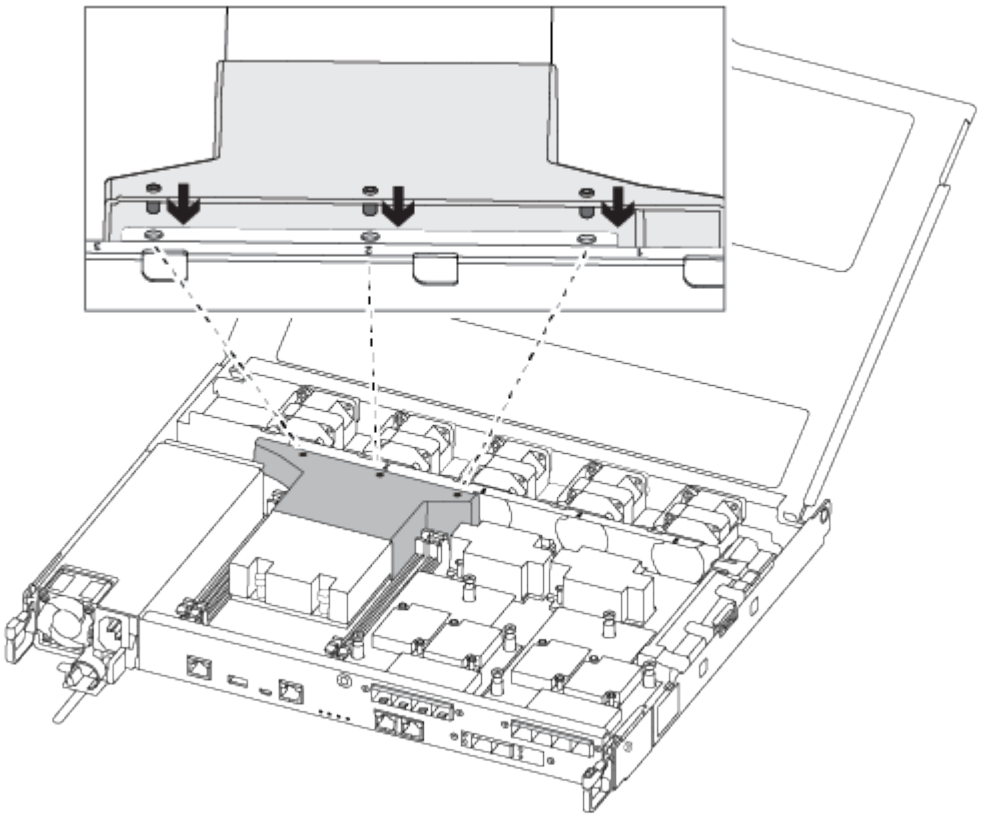

2. Close the controller module cover and tighten the thumbscrew.

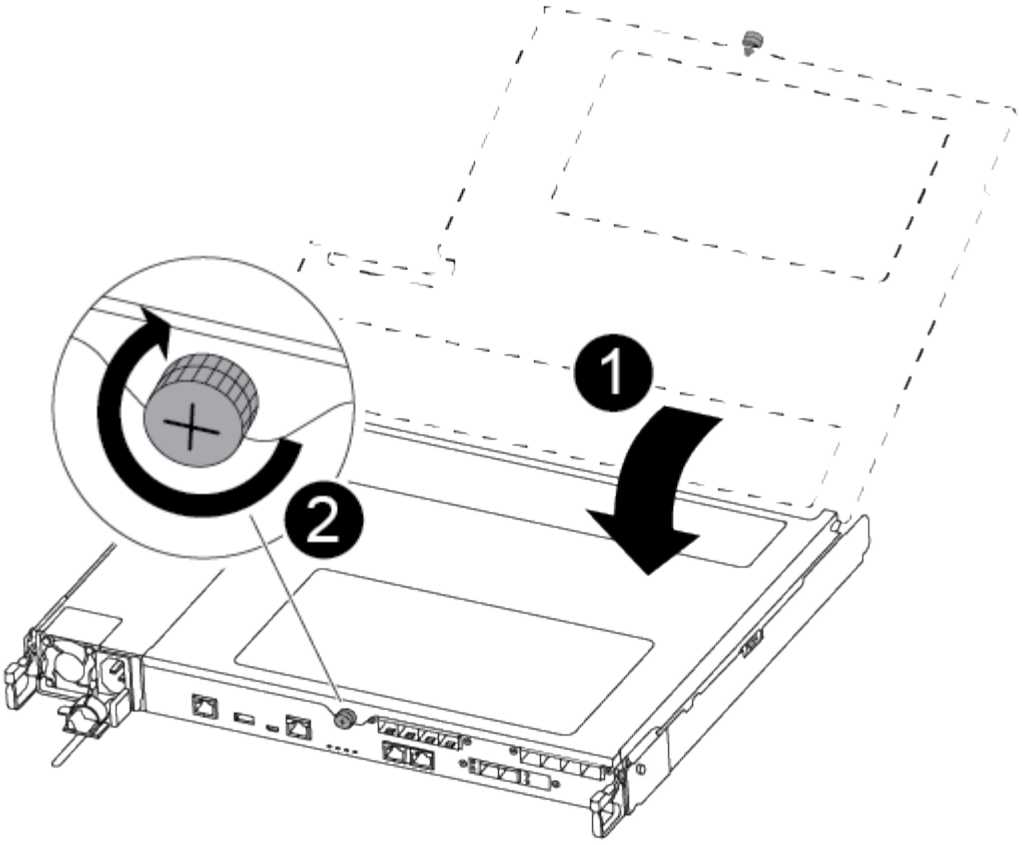

1	Controller module cover
2	Thumbscrew

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching

mechanisms snap into place.

The controller module should be fully inserted and flush with the edges of the chassis.

f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

Restore and verify the system configuration - ASA A250

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - ha
   - mcc
   - mccip
   - non-ha

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - ASA A250

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and

then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - ASA A250

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - ASA A250

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

About this task

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

> If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

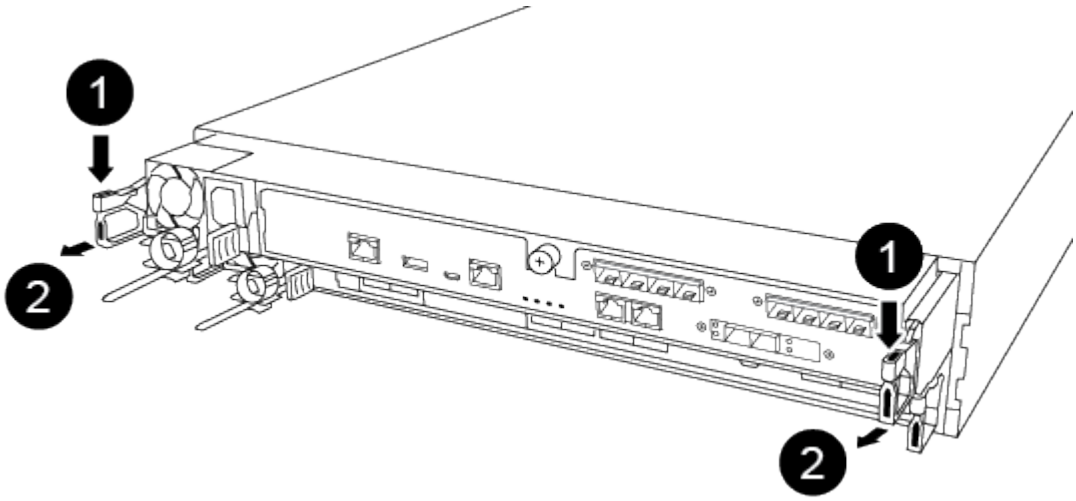

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

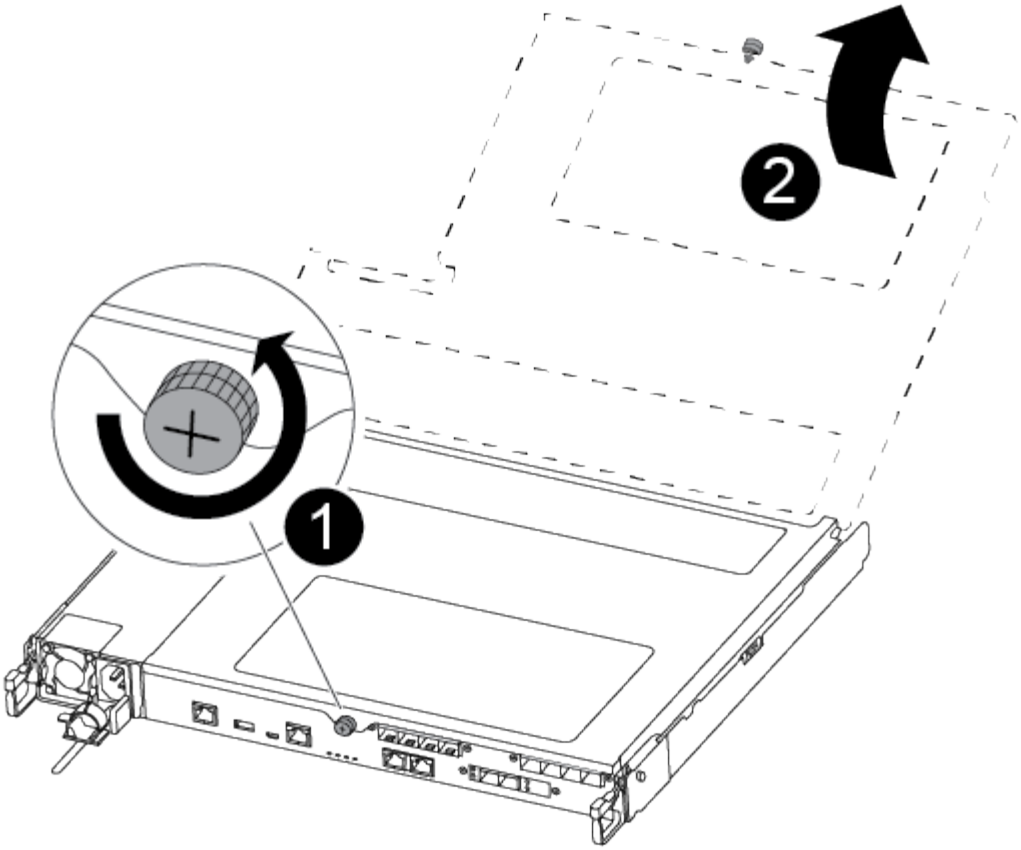

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

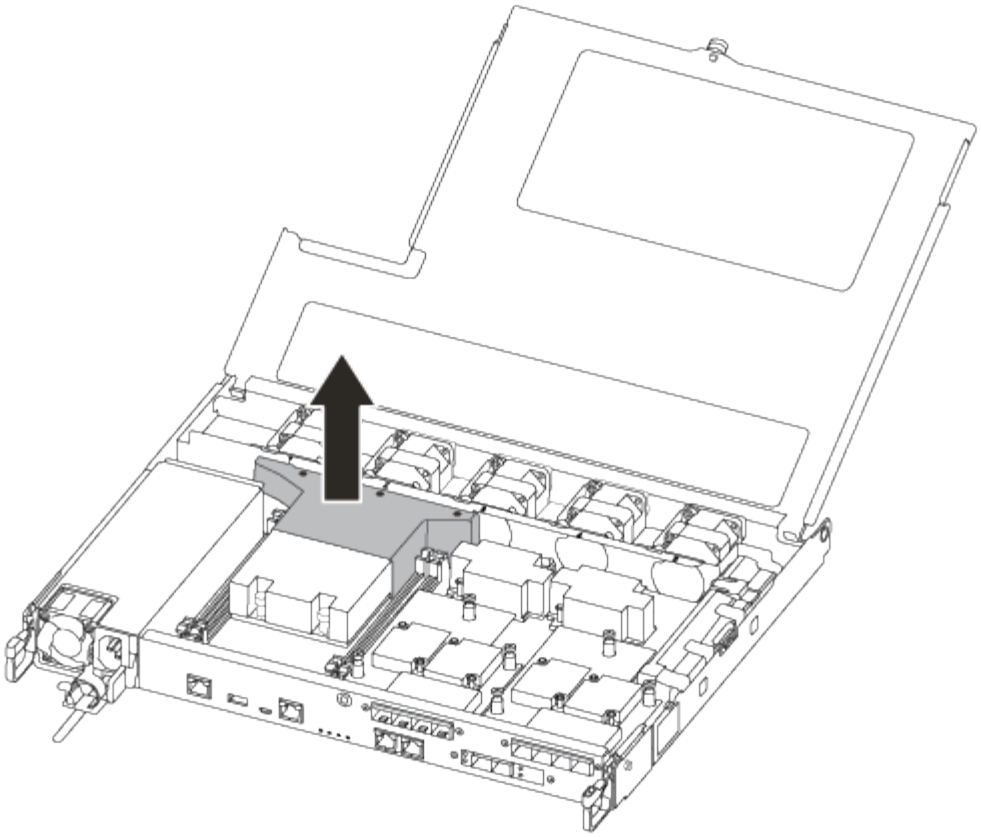

Step 3: Replace a DIMM

To replace a DIMM, you must locate it in the controller module using the DIMM map label on top of the air duct and then replace it following the specific sequence of steps.

Use the following video or the tabulated steps to replace a DIMM:

Animation - Replace a DIMM

1. Replace the impaired DIMM on your controller module.

   The DIMMs are in slot 3 or 1 on the motherboard. Slot 2 and 4 are left empty. Do not attempt to install DIMMs into these slots.

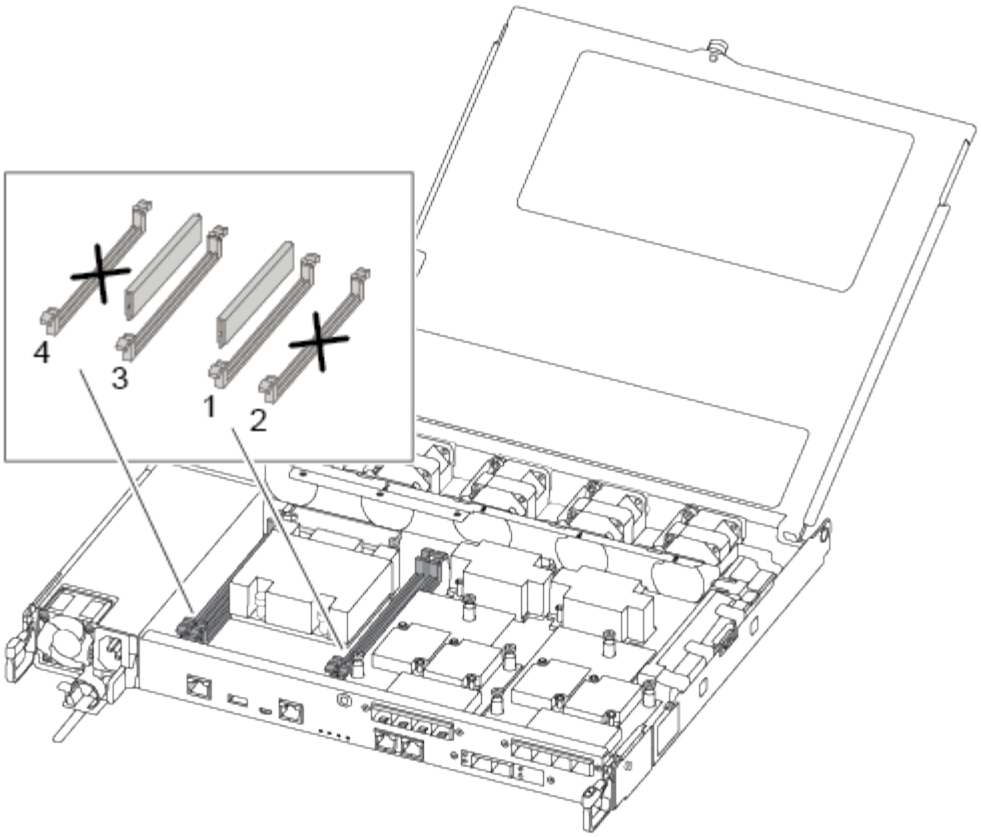

2. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
3. Slowly push apart the DIMM ejector tabs on either side of the DIMM, and slide the DIMM out of the slot.
4. Leave DIMM ejector tabs on the connector in the open position.
5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   Hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

6. Insert the replacement DIMM squarely into the slot.

   The DIMMs fit tightly in the socket. If not, reinsert the DIMM to realign it with the socket.

7. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following illustrations or the written steps to install the replacement controller module in the chassis.

1. If you have not already done so, install the air duct.

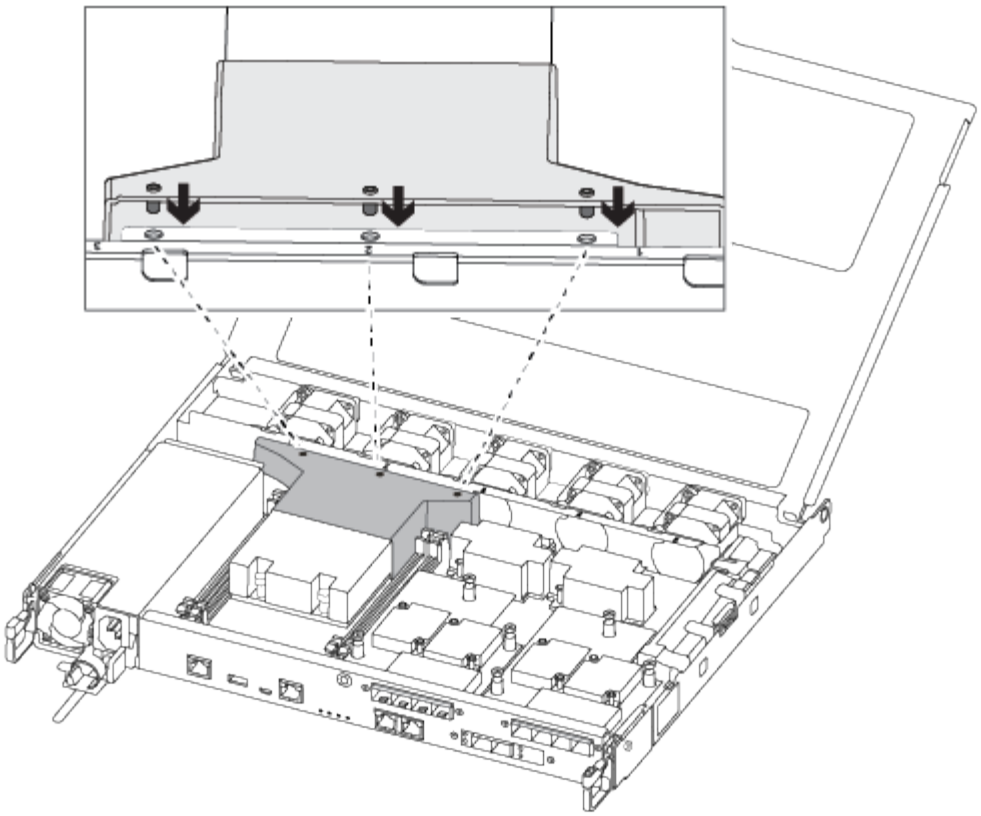

2. Close the controller module cover and tighten the thumbscrew.

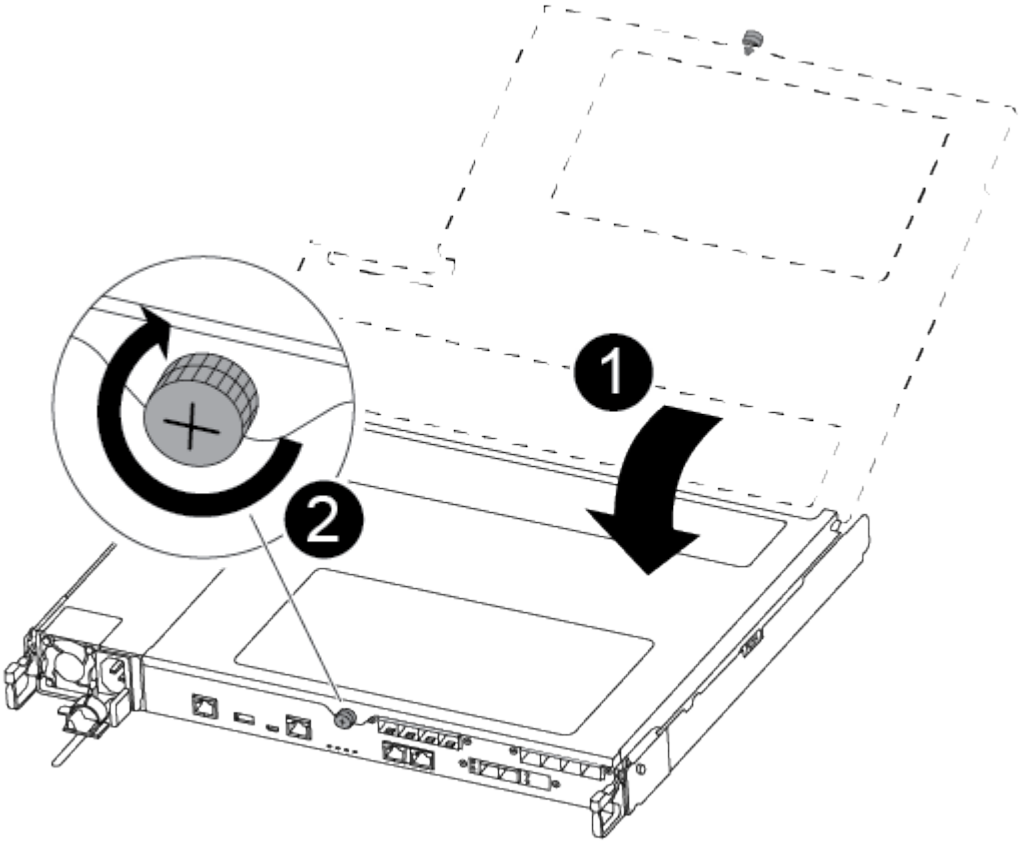

1	Controller module cover
2	Thumbscrew

3. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.
4. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.
5. Recable the system, as needed.
6. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
7. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - ASA A250

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system

console.

The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan - ASA A250

To replace a fan, remove the failed fan module and replace it with a new fan module.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

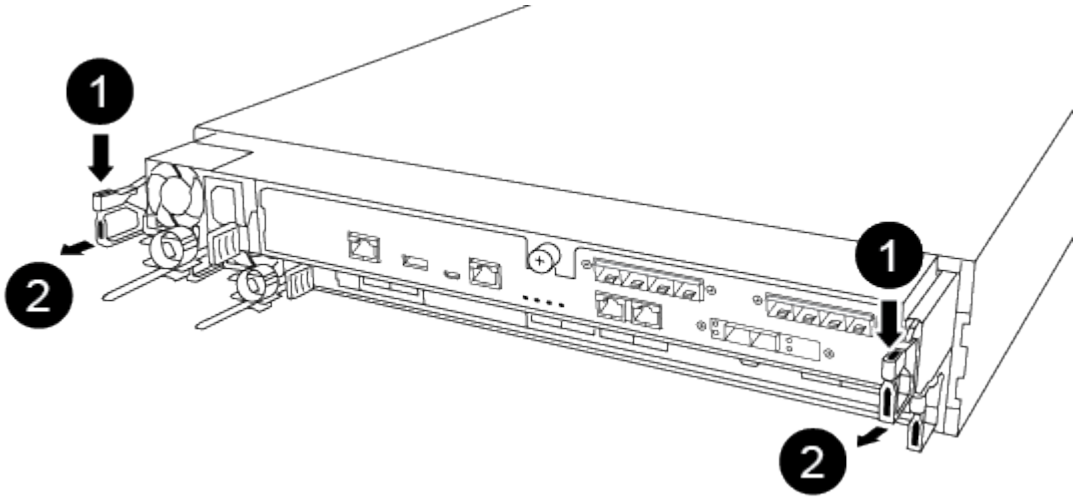

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

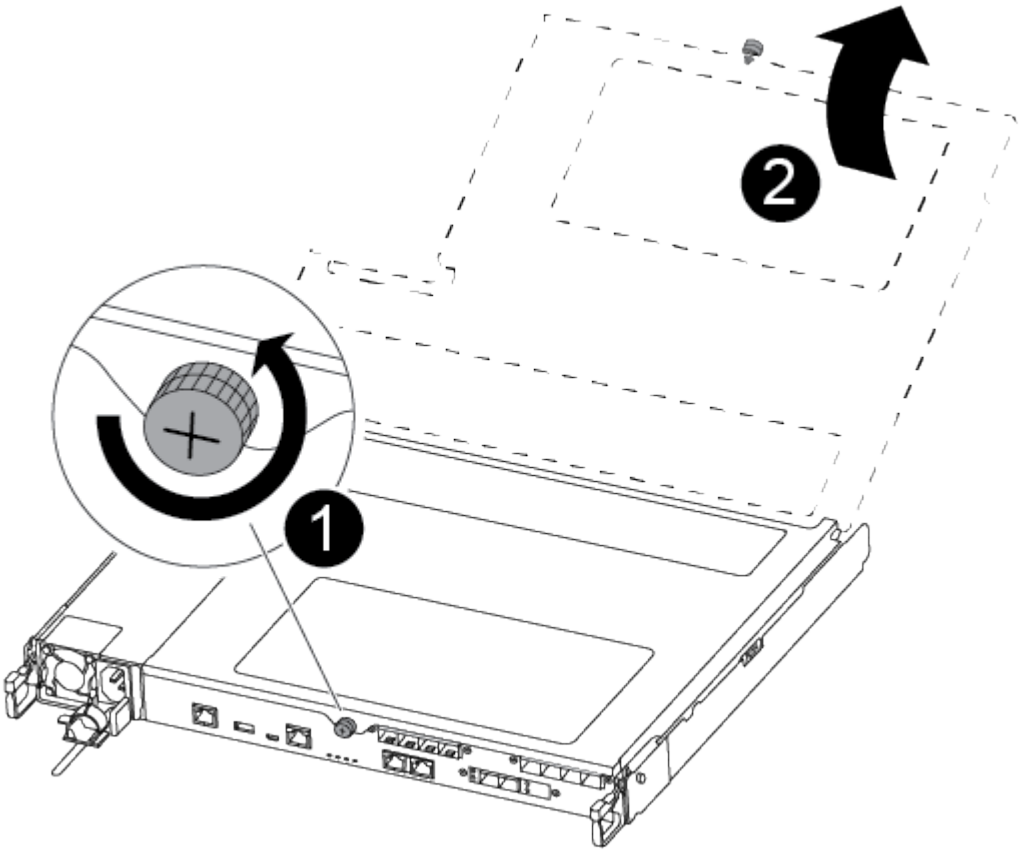

1	Thumbscrew
2	Controller module cover

Step 3: Replace a fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

Use the following video or the tabulated steps to replace a fan:

Animation - Replace a fan

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
2. Remove the fan module by pinching the side of the fan module, and then lifting the fan module straight out of the controller module.

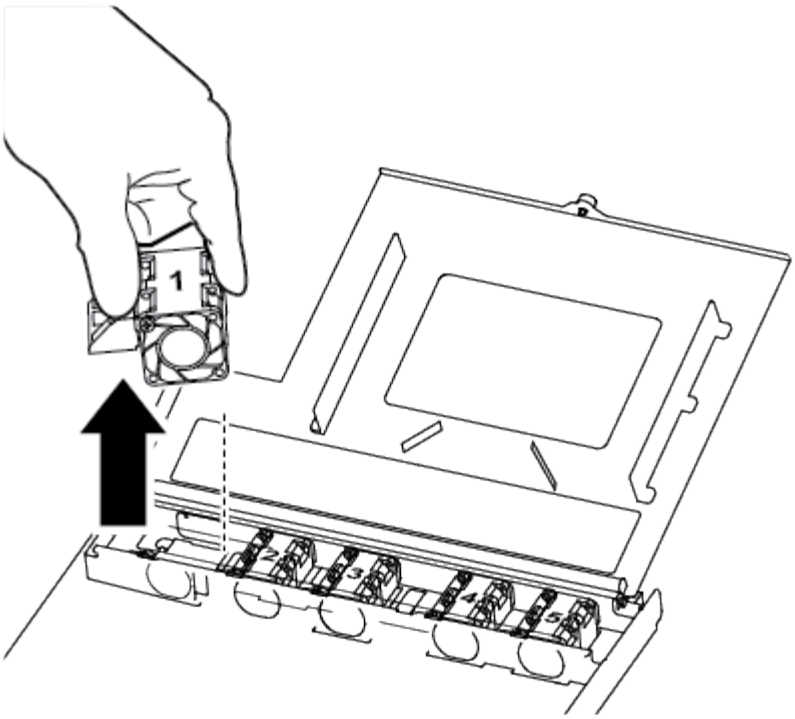

1	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Close the controller module cover and tighten the thumbscrew.

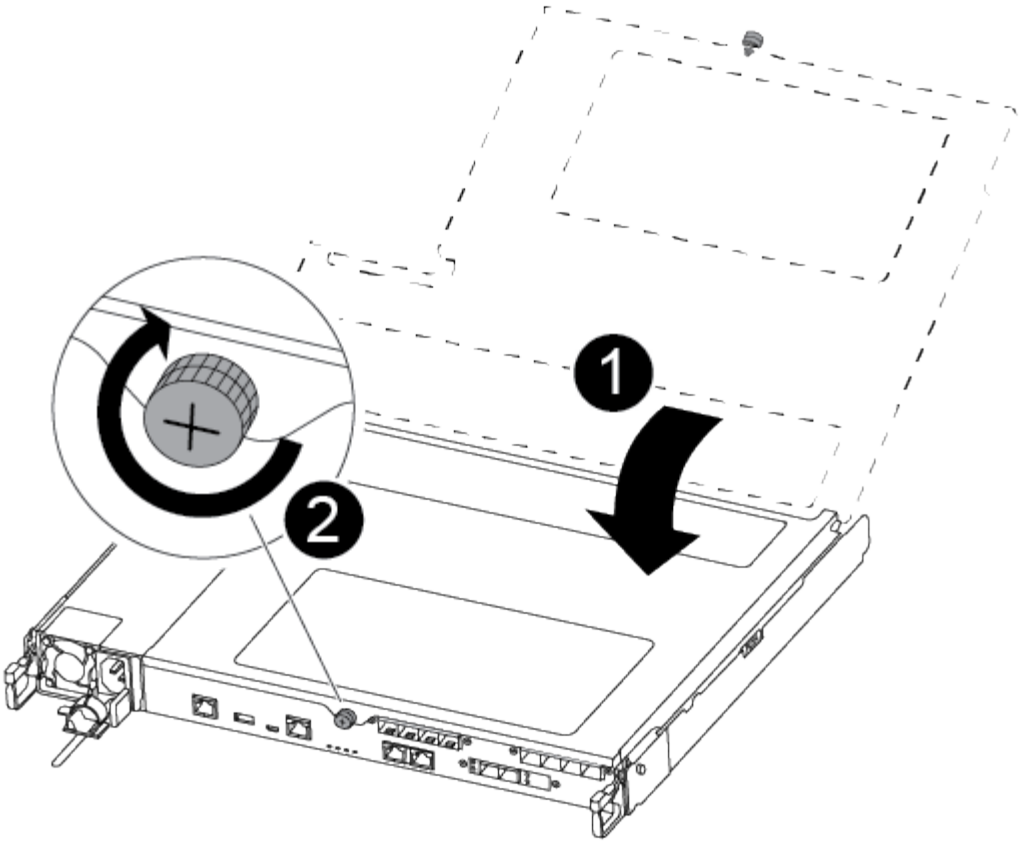

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.

   f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace or install a mezzanine card - ASA A250

To replace a failed mezzanine card, you must remove the cables and any SFP or QSFP modules, replace the card, reinstall the SFP or QSFP modules and recable the cards. To install a new mezzanine card, you must have the appropriate cables and SFP or QSFP modules.

About this task

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

Remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

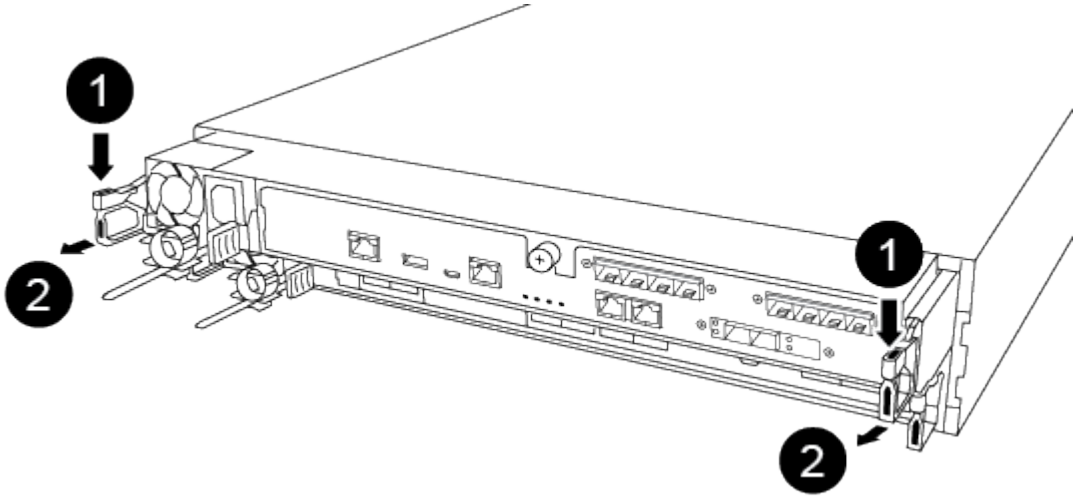

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

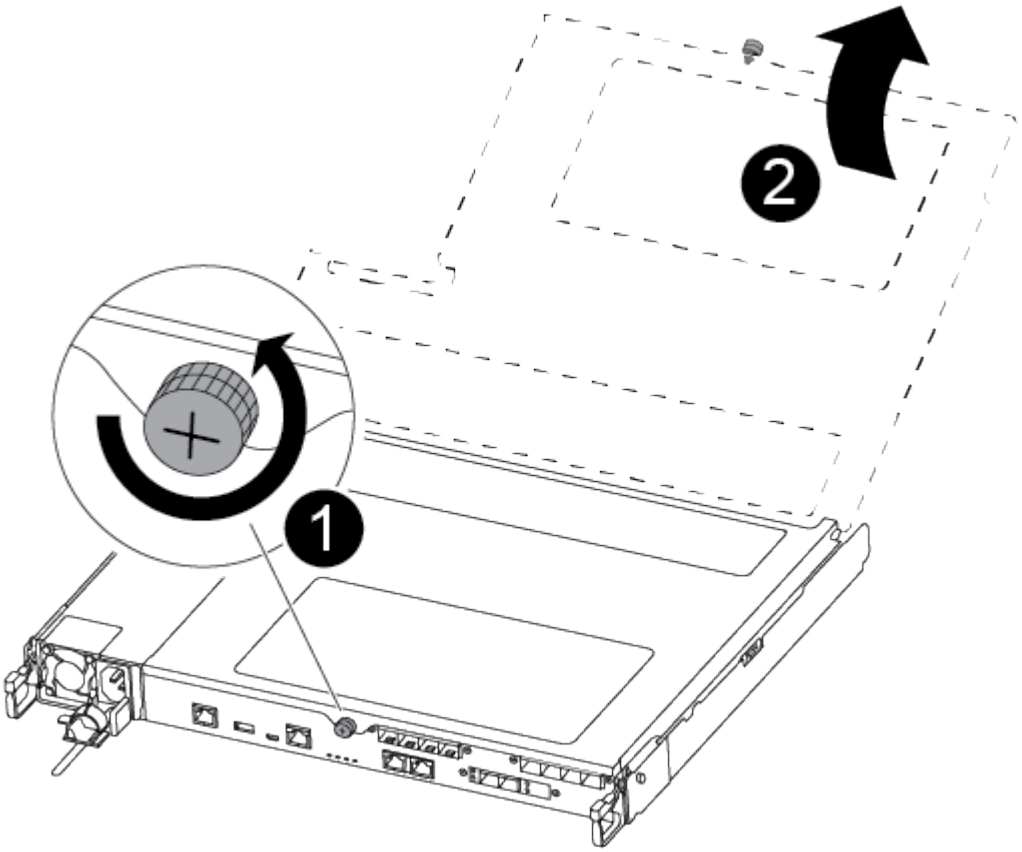

1	Thumbscrew
2	Controller module cover.

Step 3: Replace or install a mezzanine card

To replace a mezzanine card, you must remove the impaired card and install the replacement card; to install a mezzanine card, you must remove the faceplate and install the new card.

Use the following video or the tabulated steps to replace a mezzanine card:

Animation - Replace a mezzanine card

1. To replace a mezzanine card:
2. Locate and replace the impaired mezzanine card on your controller module.

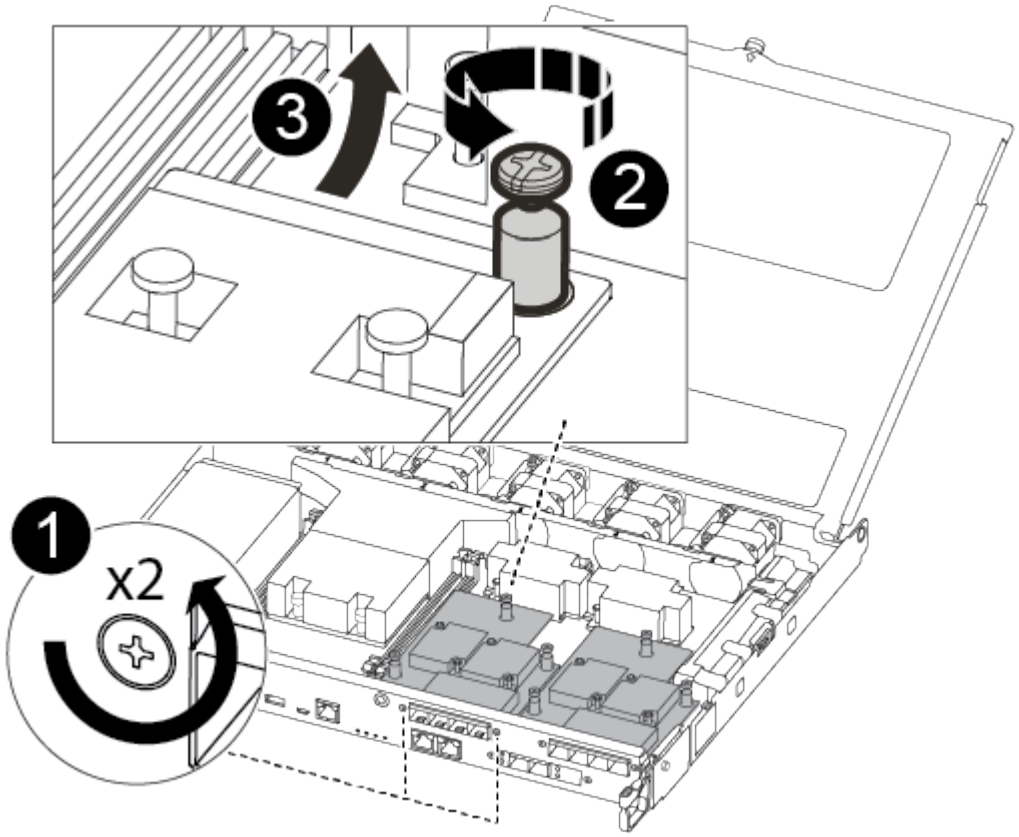

1	Remove screws on the face of the controller module.
2	Loosen the screw in the controller module.
3	Remove the mezzanine card.

a. Unplug any cabling associated with the impaired mezzanine card.

   Make sure that you label the cables so that you know where they came from.

b. Remove any SFP or QSFP modules that might be in the impaired mezzanine card and set it aside.
c. Using the #1 magnetic screwdriver, remove the screws from the face of the controller module and set them aside safely on the magnet.
d. Using the #1 magnetic screwdriver, loosen the screw on the impaired mezzanine card.
e. Using the #1 magnetic screwdriver, gently lift the impaired mezzanine card directly out of the socket and set it aside.
f. Remove the replacement mezzanine card from the antistatic shipping bag and align it to the inside face of the controller module.
g. Gently align the replacement mezzanine card into place.

h. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the controller module and on the mezzanine card.

> Do not apply force when tightening the screw on the mezzanine card; you might crack it.

i. Insert any SFP or QSFP modules that were removed from the impaired mezzanine card to the replacement mezzanine card.

3. To install a mezzanine card:
4. You install a new mezzanine card if your system does not have one.
   a. Using the #1 magnetic screwdriver, remove the screws from the face of the controller module and the faceplate covering the mezzanine card slot, and set them aside safely on the magnet.
   b. Remove the mezzanine card from the antistatic shipping bag and align it to the inside face of the controller module.
   c. Gently align the mezzanine card into place.
   d. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the controller module and on the mezzanine card.

> Do not apply force when tightening the screw on the mezzanine card; you might crack it.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Close the controller module cover and tighten the thumbscrew.

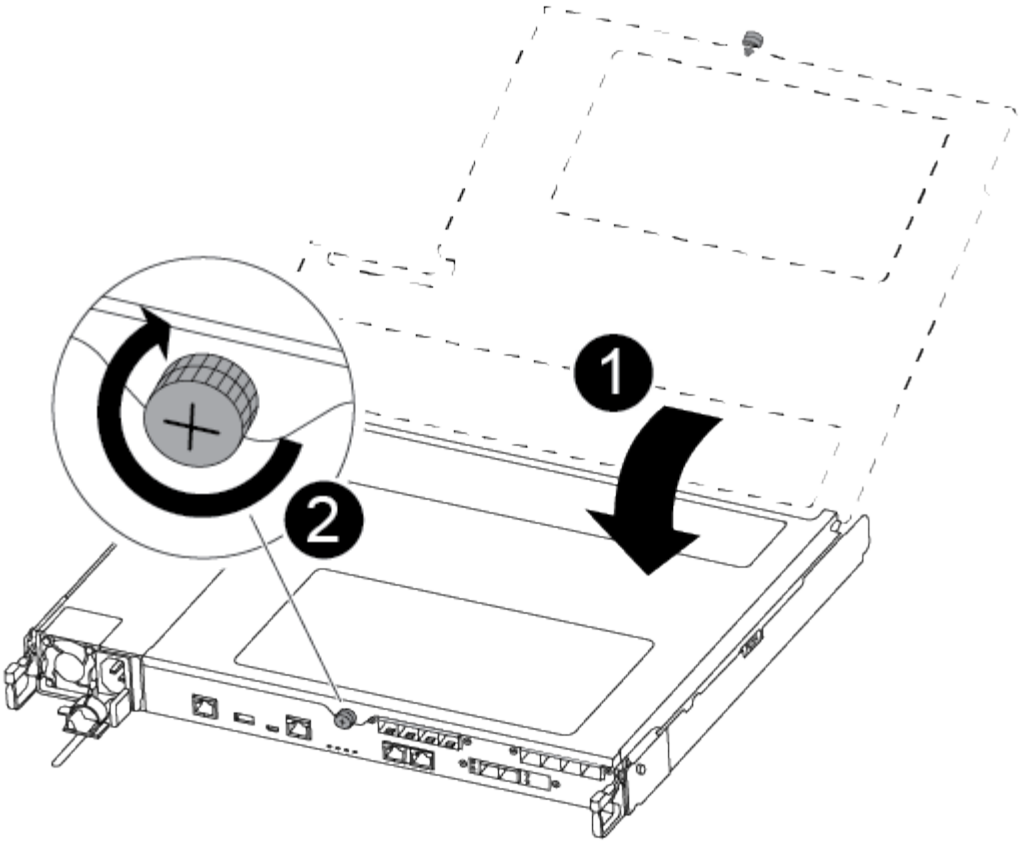

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.

   f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVMEM battery - ASA A250

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

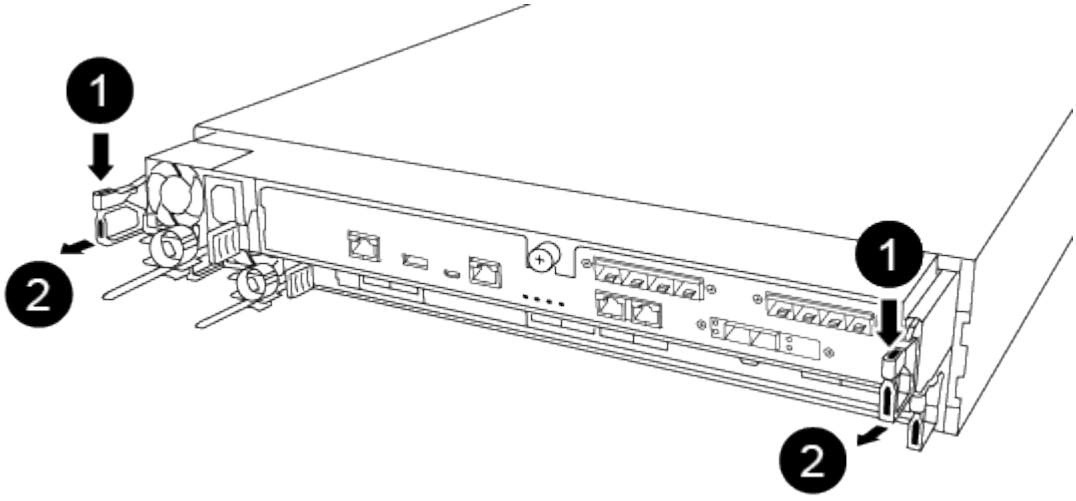

1	Lever

2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

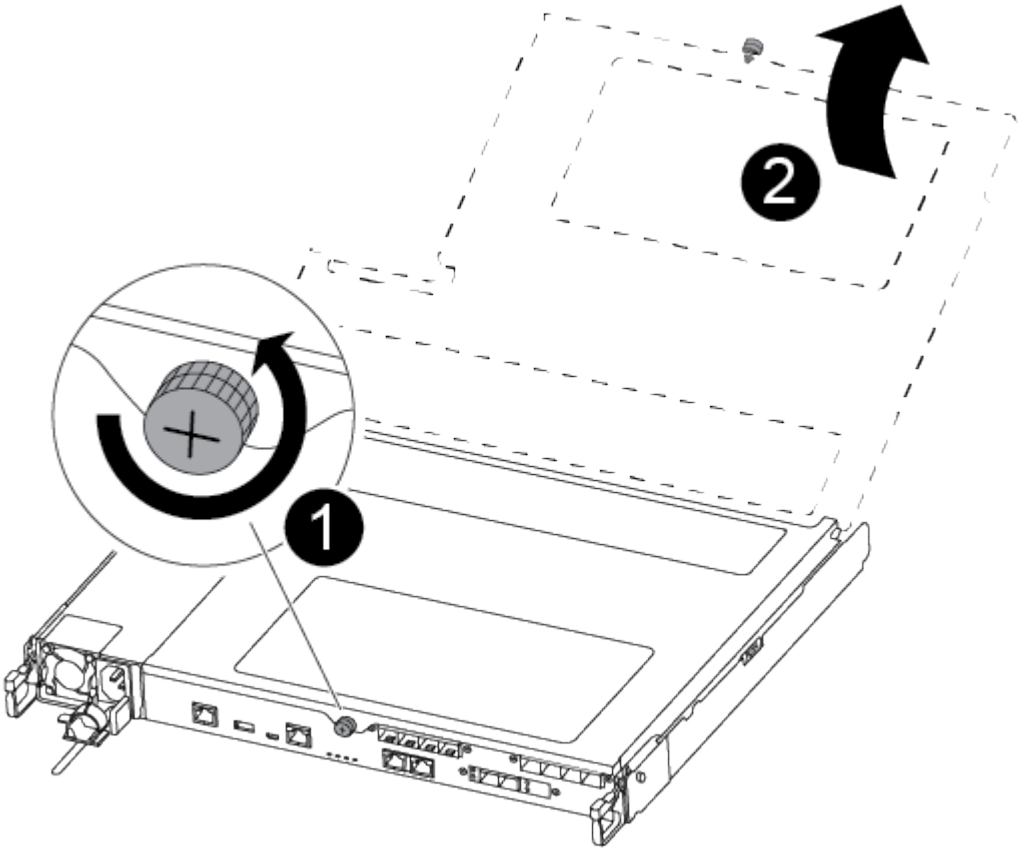

1	Thumbscrew
2	Controller module cover.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module.

Use the following video or the tabulated steps to replace the NVMEM battery:

Animation - Replace the NVMEM battery

1. Locate and replace the impaired NVMEM battery on your controller module.

It is recommended that you follow the illustrated instructions in the order listed.

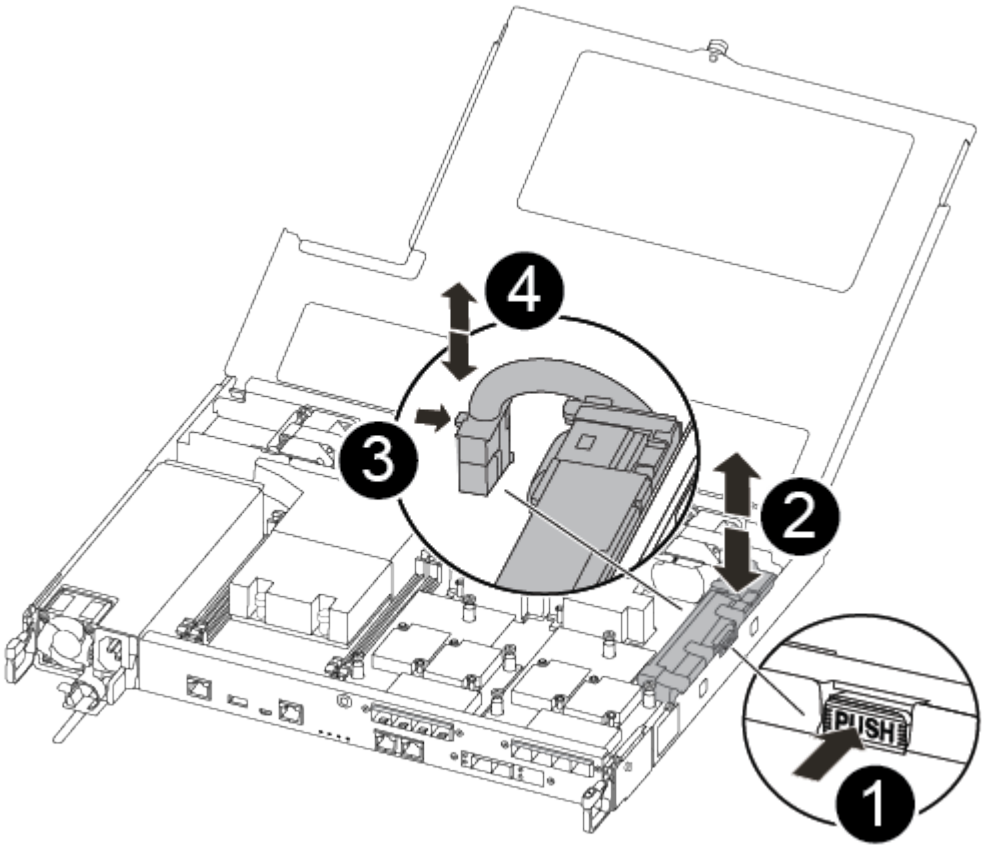

1	Squeeze the clip on the face of the battery plug.
2	Unplug the battery cable from the socket.
3	Grasp the battery and press the blue locking tab marked PUSH.
4	Lift the battery out of the holder and controller module.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module and set it aside.
4. Remove the replacement NV battery from the antistatic shipping bag and align it to the battery holder.
5. Insert the replacement NV battery plug into the socket.
6. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
7. Press firmly down on the battery pack to make sure that it is locked into place.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following illustration or the written steps to install the replacement controller module in the chassis.

1. Close the controller module cover and tighten the thumbscrew.

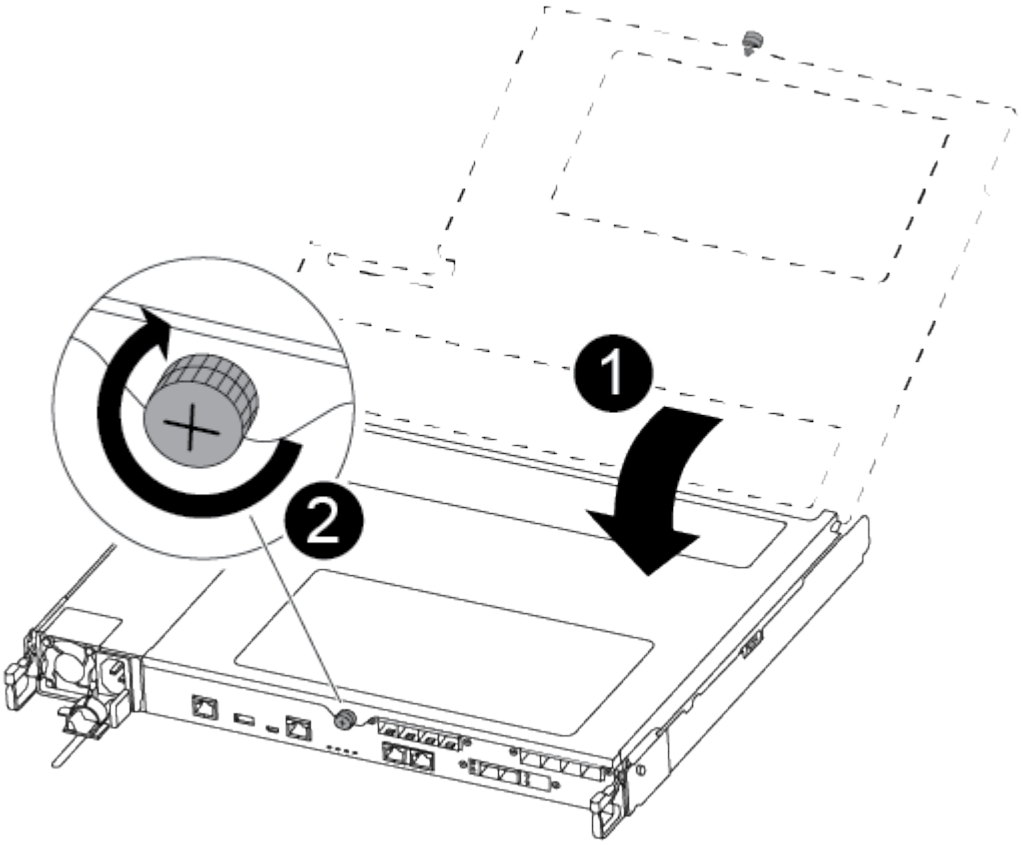

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

The controller module should be fully inserted and flush with the edges of the chassis.

f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.
4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - ASA A250

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one PSU at a time.

> It is a best practice to replace the PSU within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded PSU until the PSU is replaced.

- Power supplies are auto-ranging.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

Use the appropriate procedure for your type of PSU; AC or DC.

Option 1: Replace an AC PSU

Use the following video or the tabulated steps to replace the PSU:

Animation - Replace the AC PSU

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

> The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

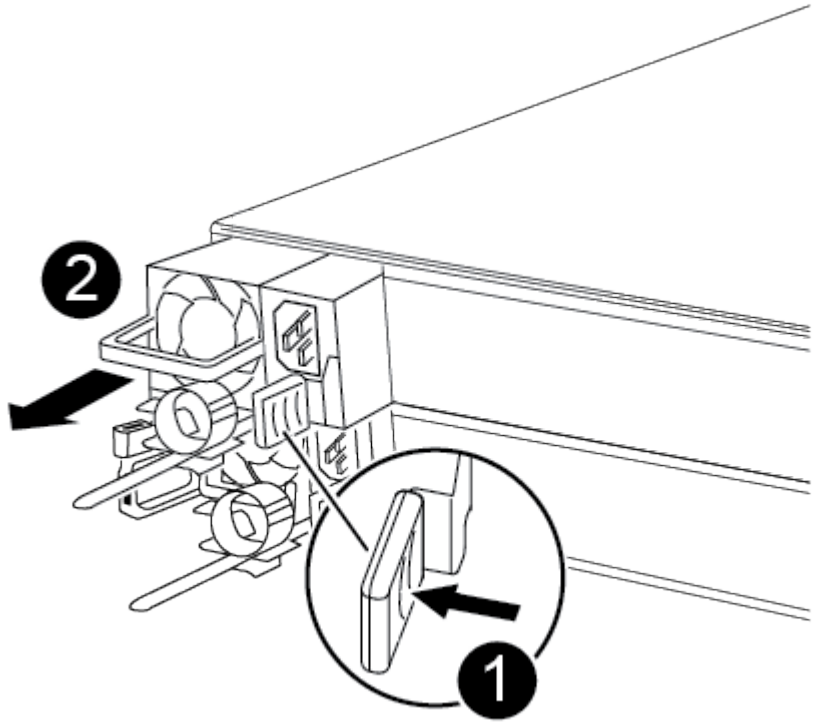

1	Blue PSU locking tab
2	Power supply

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller module.

b. Gently push the PSU into the controller module until the locking tab clicks into place.

The power supplies will only properly engage with the internal connector and lock in place one way.

> To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:

   a. Reconnect the power cable to the PSU.

   b. Secure the power cable to the PSU using the power cable retainer.

   Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC power cable connector using the thumb screws on the plug.
   b. Unplug the power cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

> The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

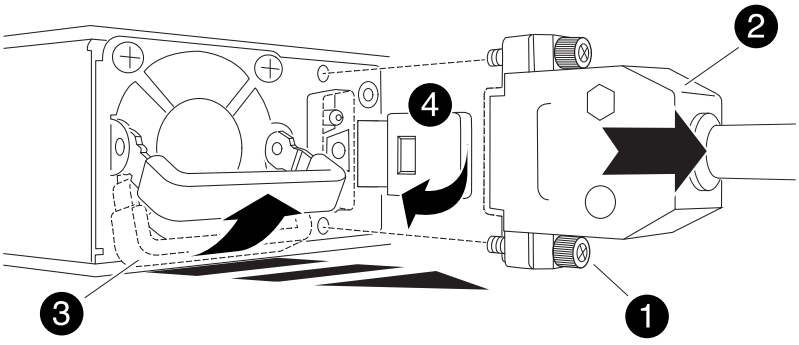

1	Thumb screws
2	D-SUB DC power cable connector

3	Power supply handle
4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA A250

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

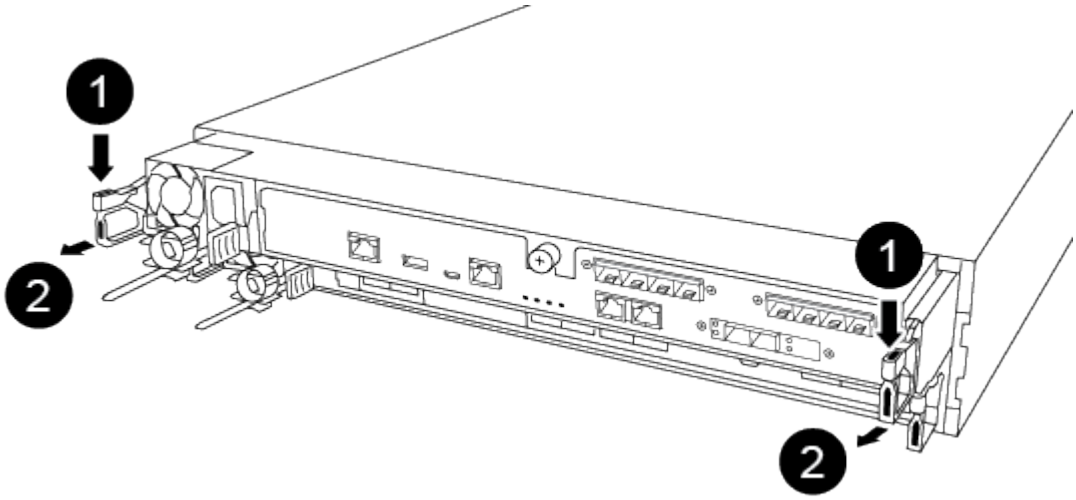

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

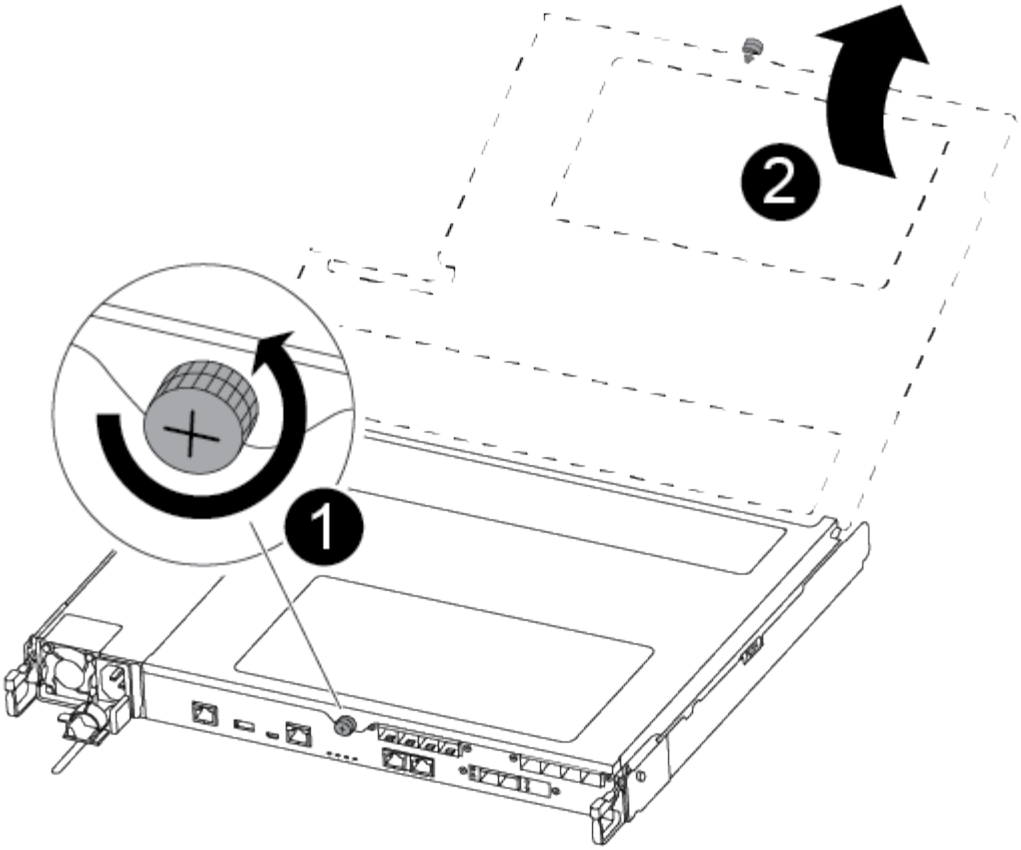

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

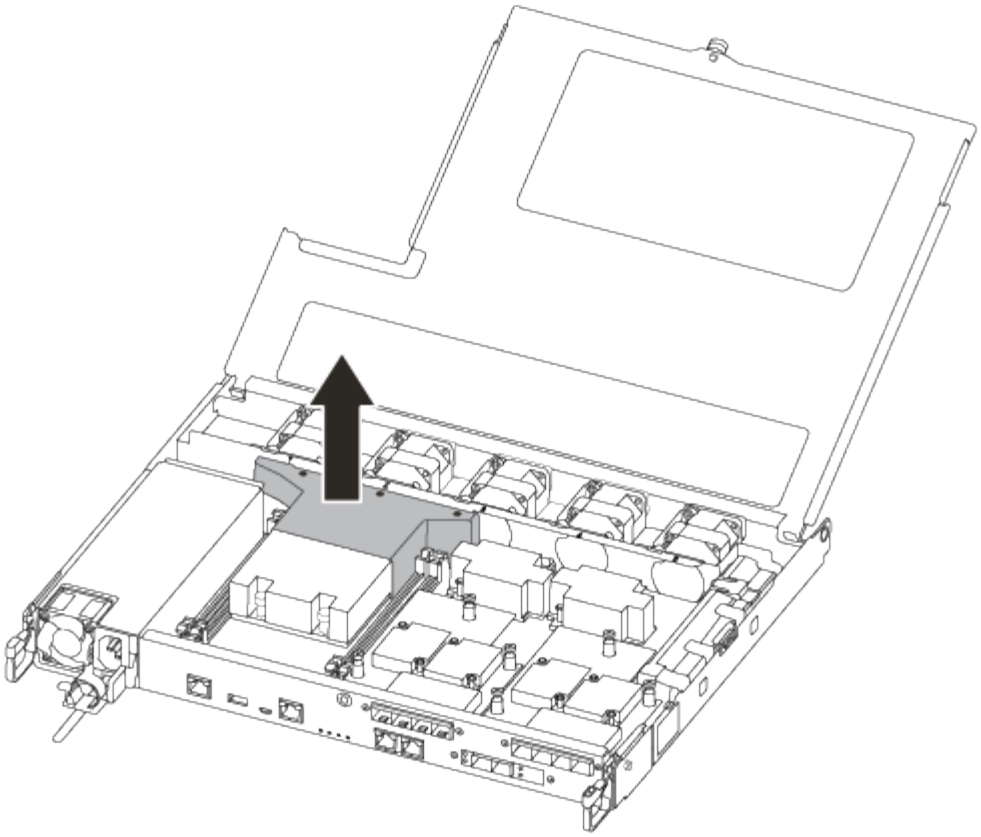

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

Use the following video or the tabulated steps to replace the RTC battery:

Animation - Replace the RTC battery

1. Locate the RTC battery between the heatsink and the midplane and remove it exactly as shown in the graphic.

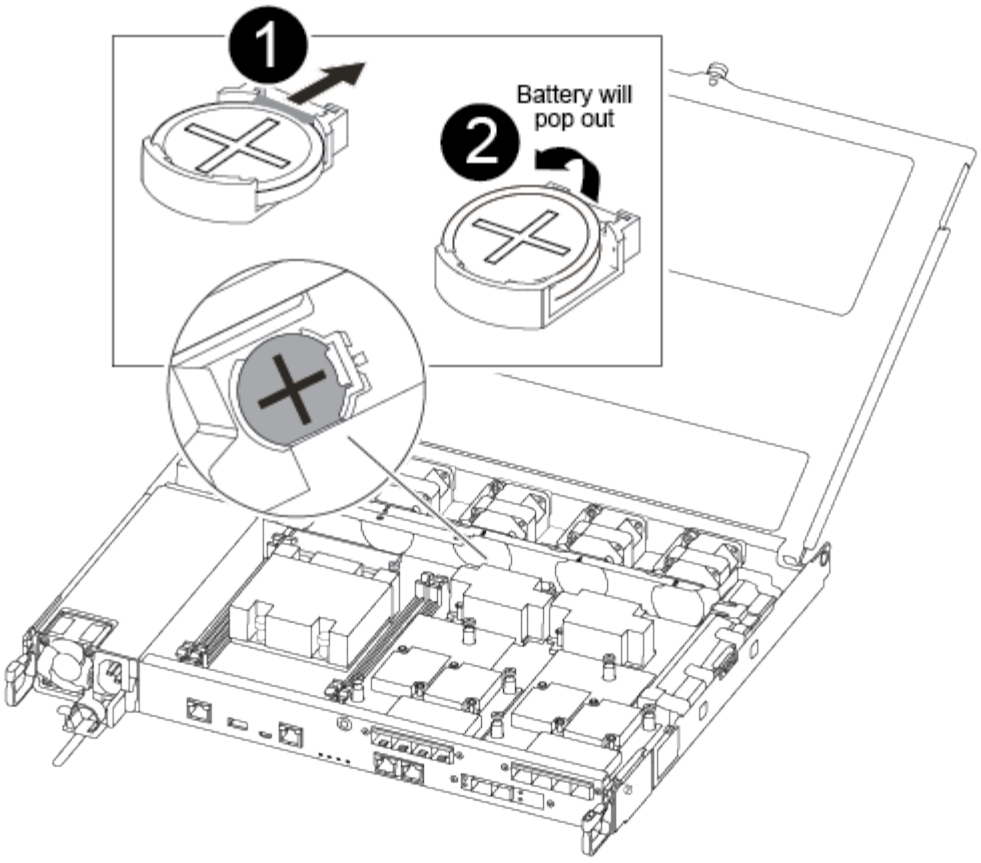

1	Gently pull tab away from the battery housing. **Attention:** Pulling it away aggressively might displace the tab.
2	Lift the battery up. **Note:** Make a note of the polarity of the battery.
3	The battery should eject out.

2. Remove the replacement battery from the antistatic shipping bag.
3. Locate the RTC battery holder between the heatsink and the midplane and insert it exactly as shown in the graphic.

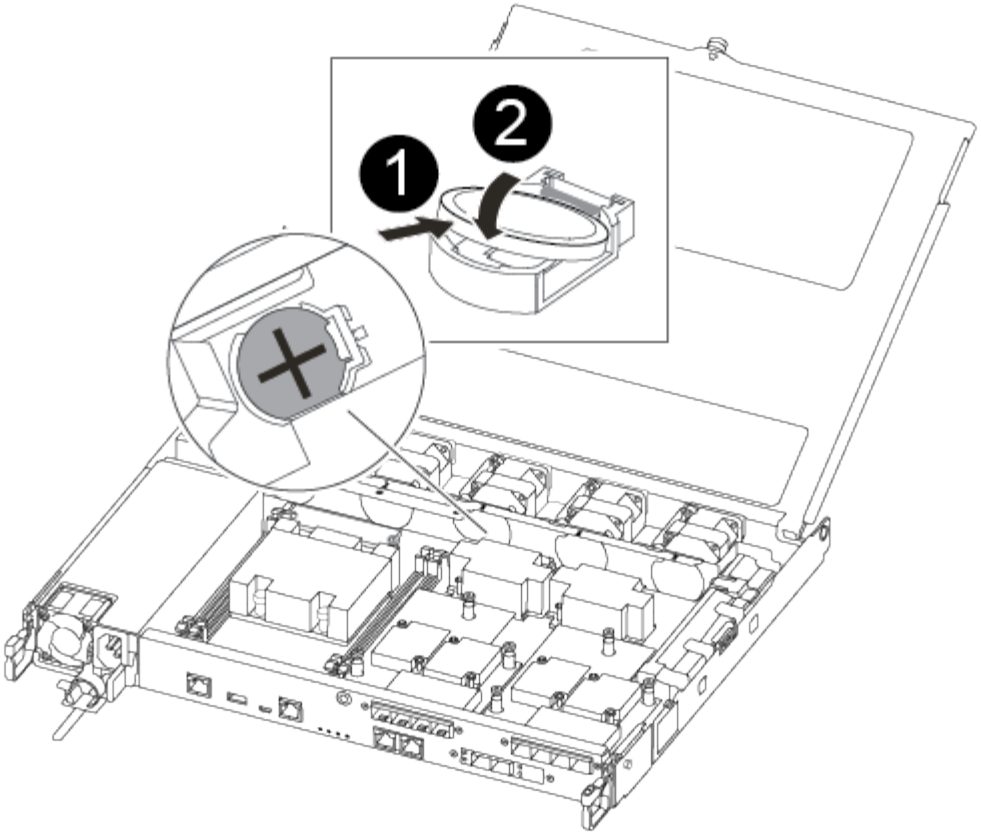

1	With positive polarity face up, slide the battery under the tab of the battery housing.
2	Push the battery gently into place and make sure the tab secures it to the housing. Pushing it in aggressively might cause the battery to eject out again.

4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Insert the controller module into the chassis:
    a. Ensure the latching mechanism arms are locked in the fully extended position.
    b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
    c. Place your index fingers through the finger holes from the inside of the latching mechanism.
    d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
    e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

       The controller module should be fully inserted and flush with the edges of the chassis.
    f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

       The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.
    g. Halt the controller at the LOADER prompt.
5. Reset the time and date on the controller:
    a. Check the date and time on the healthy controller with the `show date` command.
    b. At the LOADER prompt on the target controller, check the time and date.
    c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
    d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
    e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA A400 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

For MetroCluster configurations, see either:

- Install MetroCluster IP configuration
- Install MetroCluster Fabric-Attached configuration

Quick guide - ASA A400

The Installation and Setup instructions give graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Use the links: AFF A400 Installation and Setup Instructions.

> The ASA A400 uses the same installation procedure as the AFF A400 system.

Video steps - ASA A400

The following video shows how to install and cable your new system.

Animation - AFF A400 Installation and setup instructions

> The ASA A400 uses the same installation procedure as the AFF A400 system.

Detailed guide - ASA A400

This page provides detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your system, you need to create an account, register the system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
100 GbE cable (QSF(28)	X66211A-05 (112-00595), 0.5m X66211A-1 (112-00573), 1m X66211A-2 (112-00574), 2m X66211A-5 (112-00574), 5m		Storage, cluster interconnect/HA, and Ethernet data (order-dependent)
25 GbE cable (SFP28s)	X66240-2 (112-00598), 2m X66240-5 (112-00639), 5m		GbE network connection (order-dependent)
32 Gb FC (SFP+ Op)	X66250-2 (112-00342), 2m X66250-5 (112-00344), 5m X66250-15 (112-00346), 15m		FC network connection
Storage Cables	X66030A (112-00435), .5m X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		mini-SAS HD to mini-SAS HD cables (order-dependent)

Type of cable…	Part number and length	Connector type	For…
Optical cables	X66250-2-N-C (112-00342)		16 Gb FC or 25GbE cables for mezzanine cards (order-dependent)
RJ-45 (order dependent)	X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network
Micro-USB console cable	Not applicable		Console connection used during software setup if laptop or console does not support network discovery.
Power cables	Not applicable		Powering up the system

4. Review the *NetApp ONTAP Configuration Guide* and collect the required information listed in that guide.

   ONTAP Configuration Guide

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

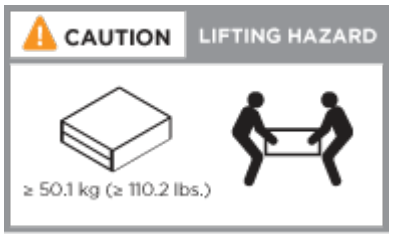

3. Attach cable management devices (as shown).

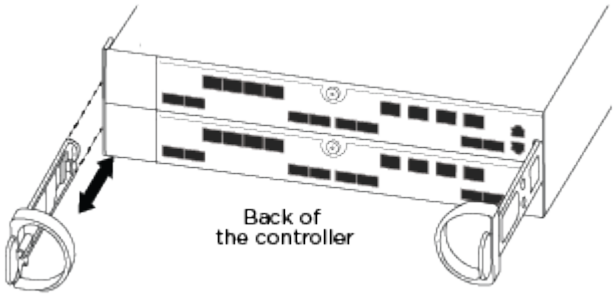

4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

If the port labels on the card are not visible, check the card installation orientation (the PCIe connector socket is on the left side of the card slot in the A400 and FAS8300/8700), and then look for the card, by part number, in the NetApp Hardware Universe for a graphic of the bezel which will show the port labels. The card part number can be found using the `sysconfig -a` command or on the system packing list.

If you are cabling an MetroCluster IP configuration, ports e0a/e0b are available for hosting data LIFs (usually in Default IPSpace).

Option 1: Cable a two-node switchless cluster

The optional data ports, optional NIC cards, and management ports on the controller modules are connected to switches. The cluster interconnect and HA ports are cabled on both controller modules.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.

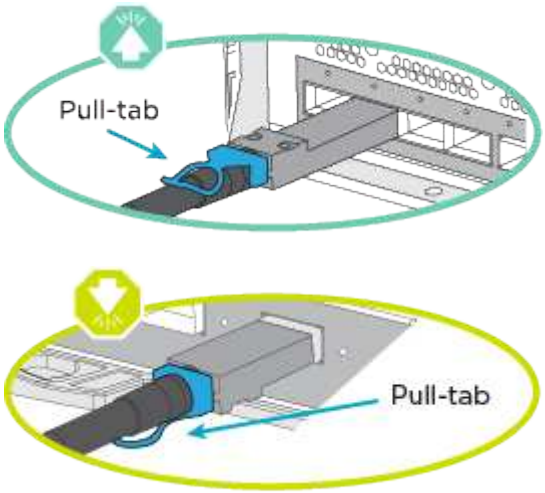

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Two-node switchless cluster cabling

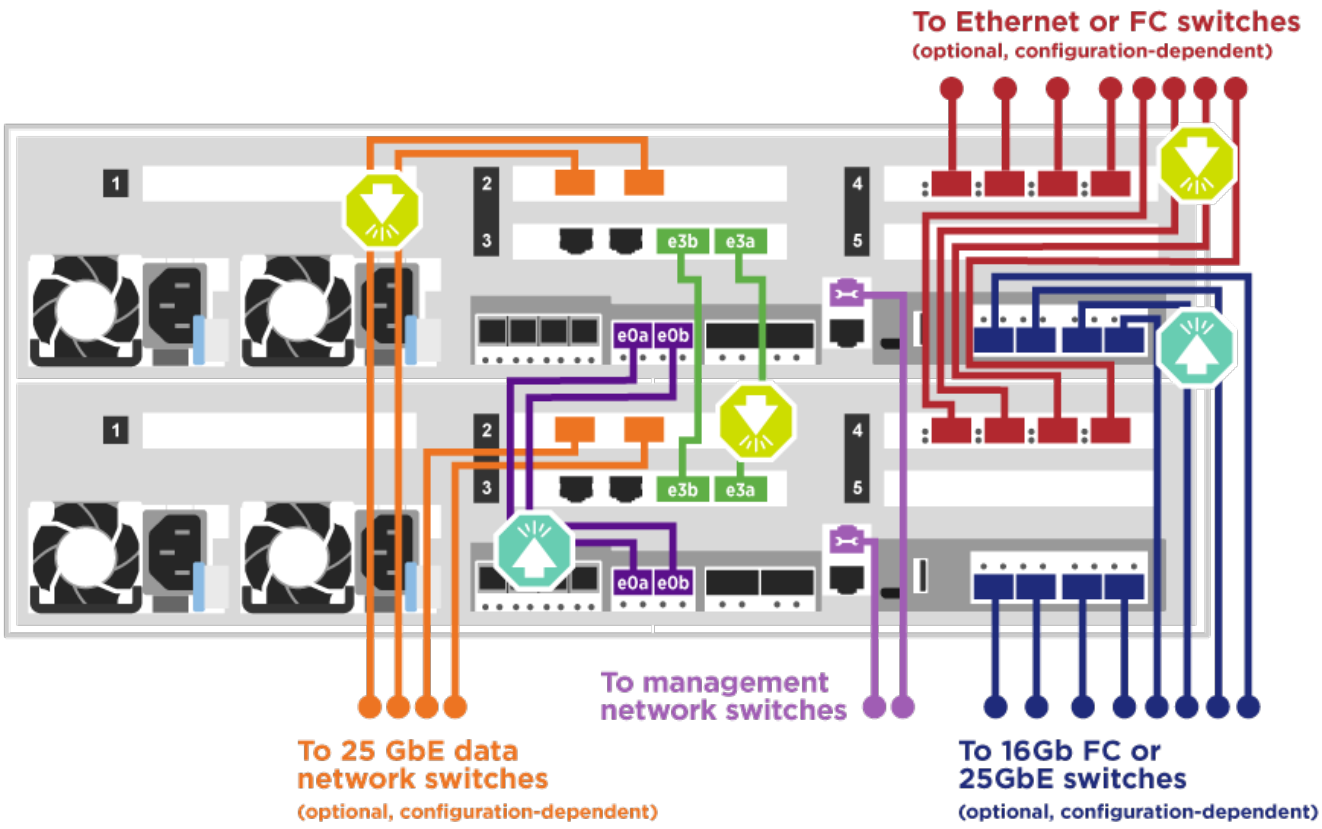

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Option 2: Cable a switched cluster

The optional data ports, optional NIC cards, mezzanine cards, and management ports on the controller modules are connected to switches. The cluster interconnect and HA ports are cabled on to the cluster/HA switch.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.

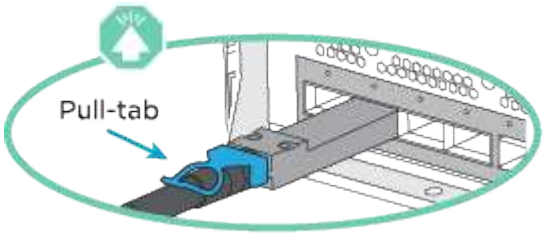

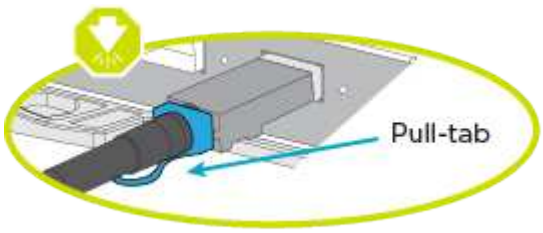

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Switched cluster cabling

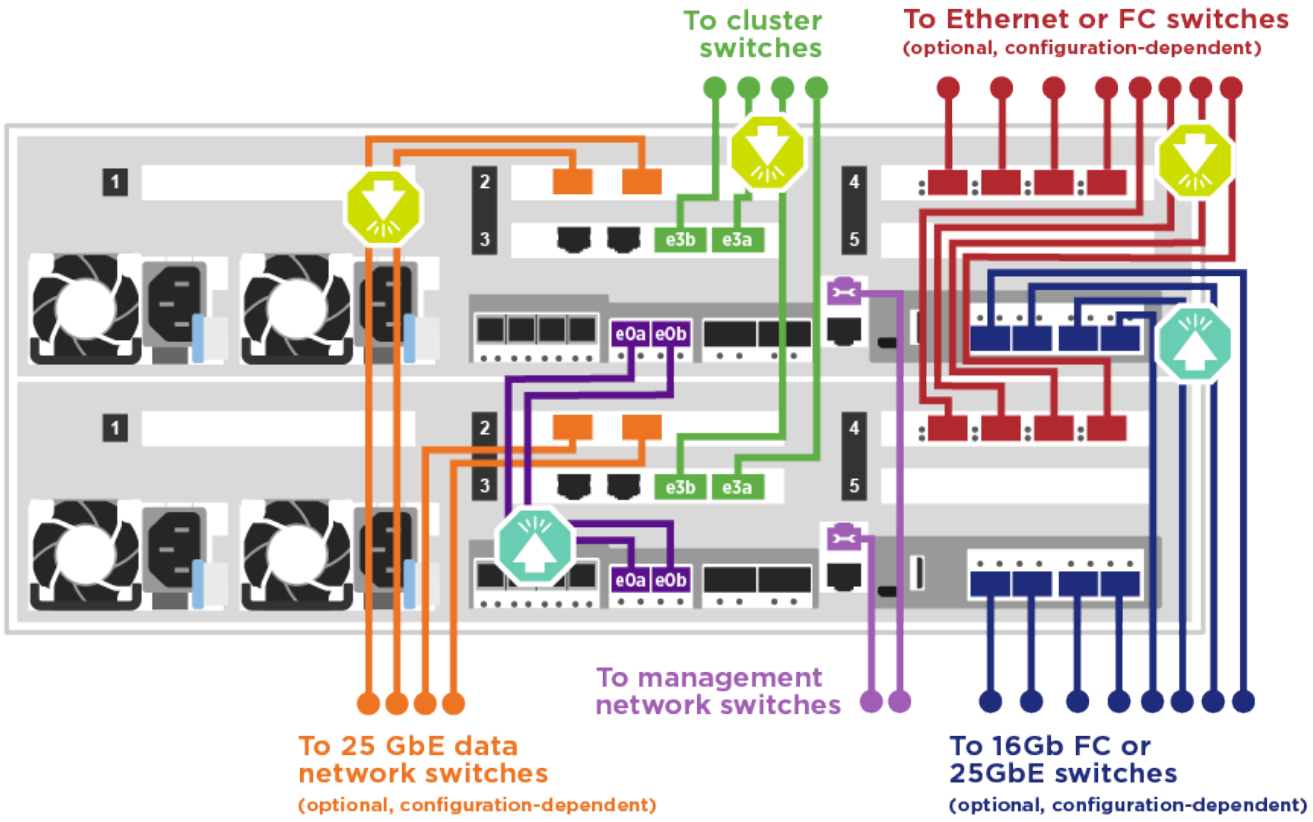

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Step 4: Cable controllers to drive shelves

You can cable either NSS224 or SAS shelves to you system.

Option 1: Cable the controllers to a single drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the NS224 are up.

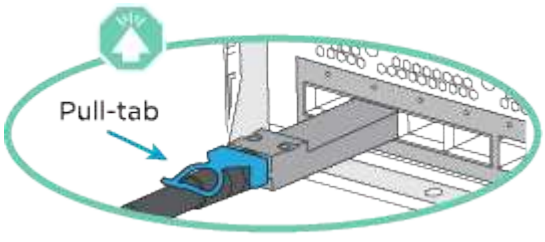

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following animation or illustration to cable your controllers to a single drive shelf.

   Animation - Cable the controllers to one NS224 drive shelf

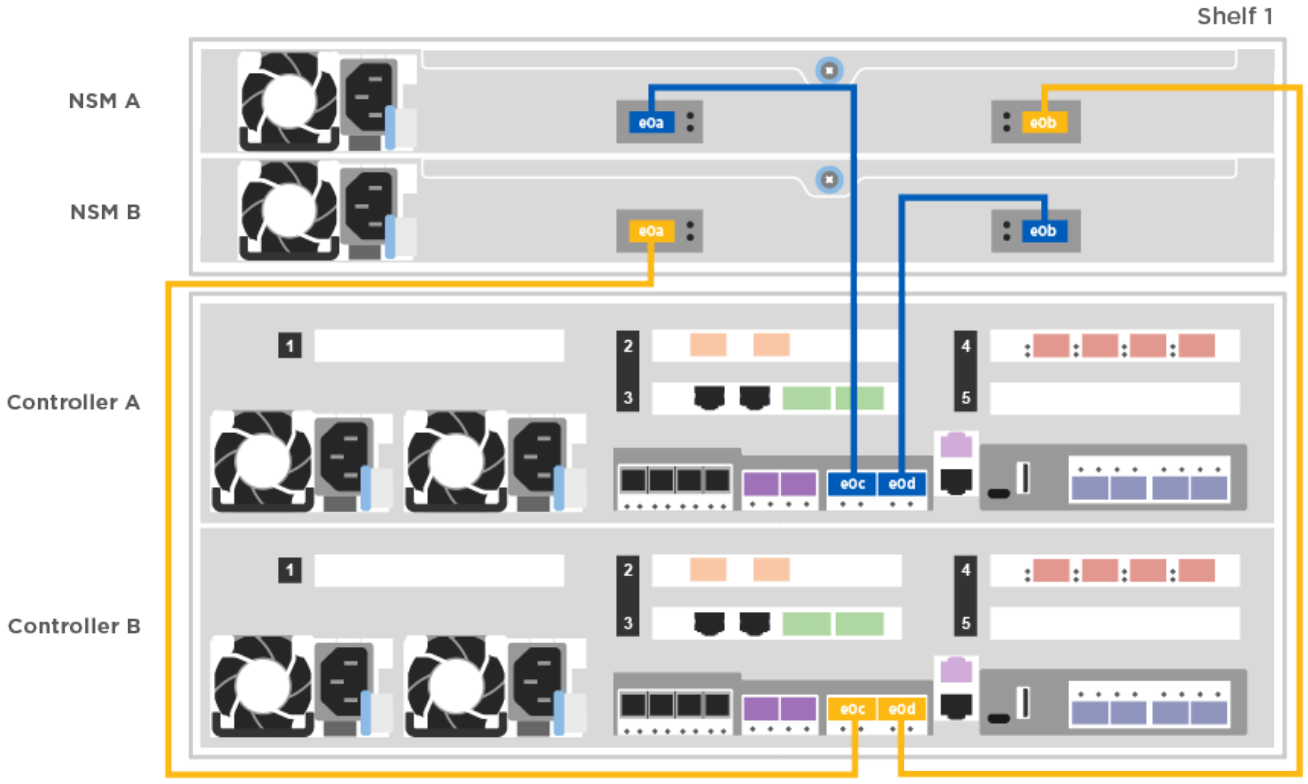

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Option 2: Cable the controllers to two drive shelves

You must cable each controller to the NSM modules on both NS224 drive shelves.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the NS224 are up.

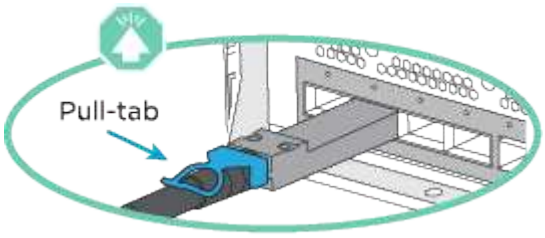

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following animation or illustration to cable your controllers to two drive shelves.

   Animation - Cable the controllers to one NS224 drive shelf

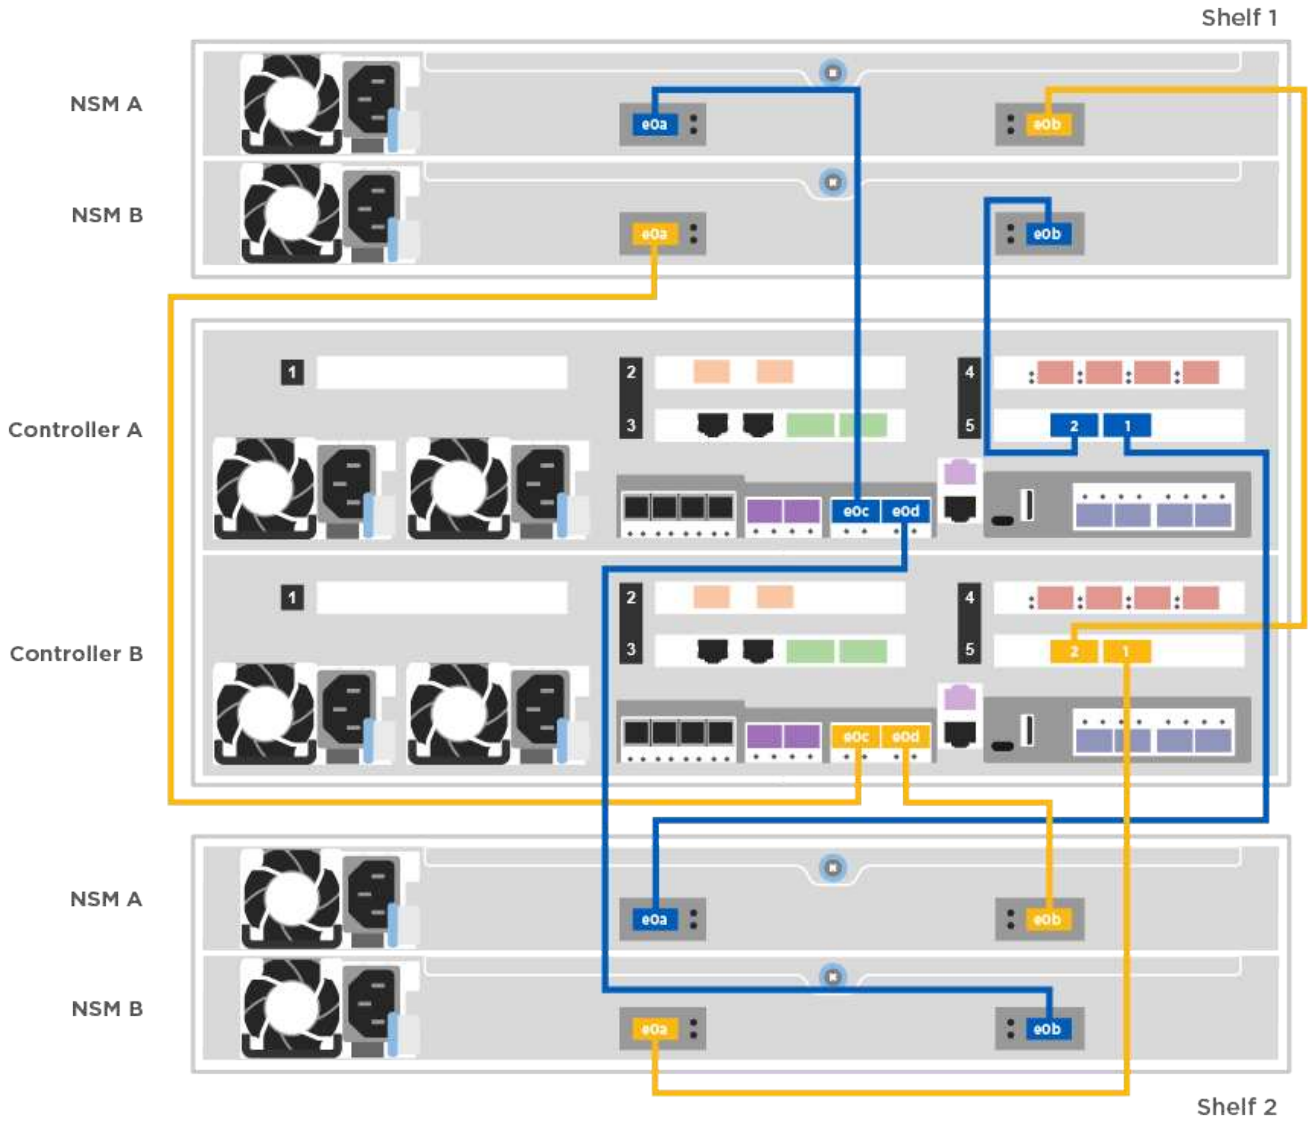

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Option 3: Cable the controllers to SAS drive shelves

You must cable each controller to the IOM modules on both SAS drive shelves.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the DS224-C are down.

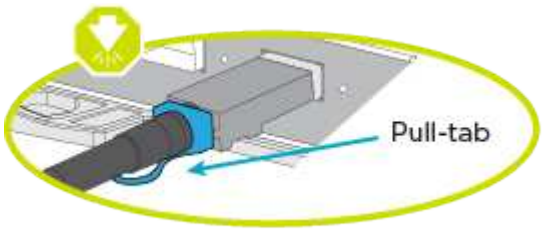

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following illustration to cable your controllers to two drive shelves.

   Animation - Cable the controllers to SAS drive shelves

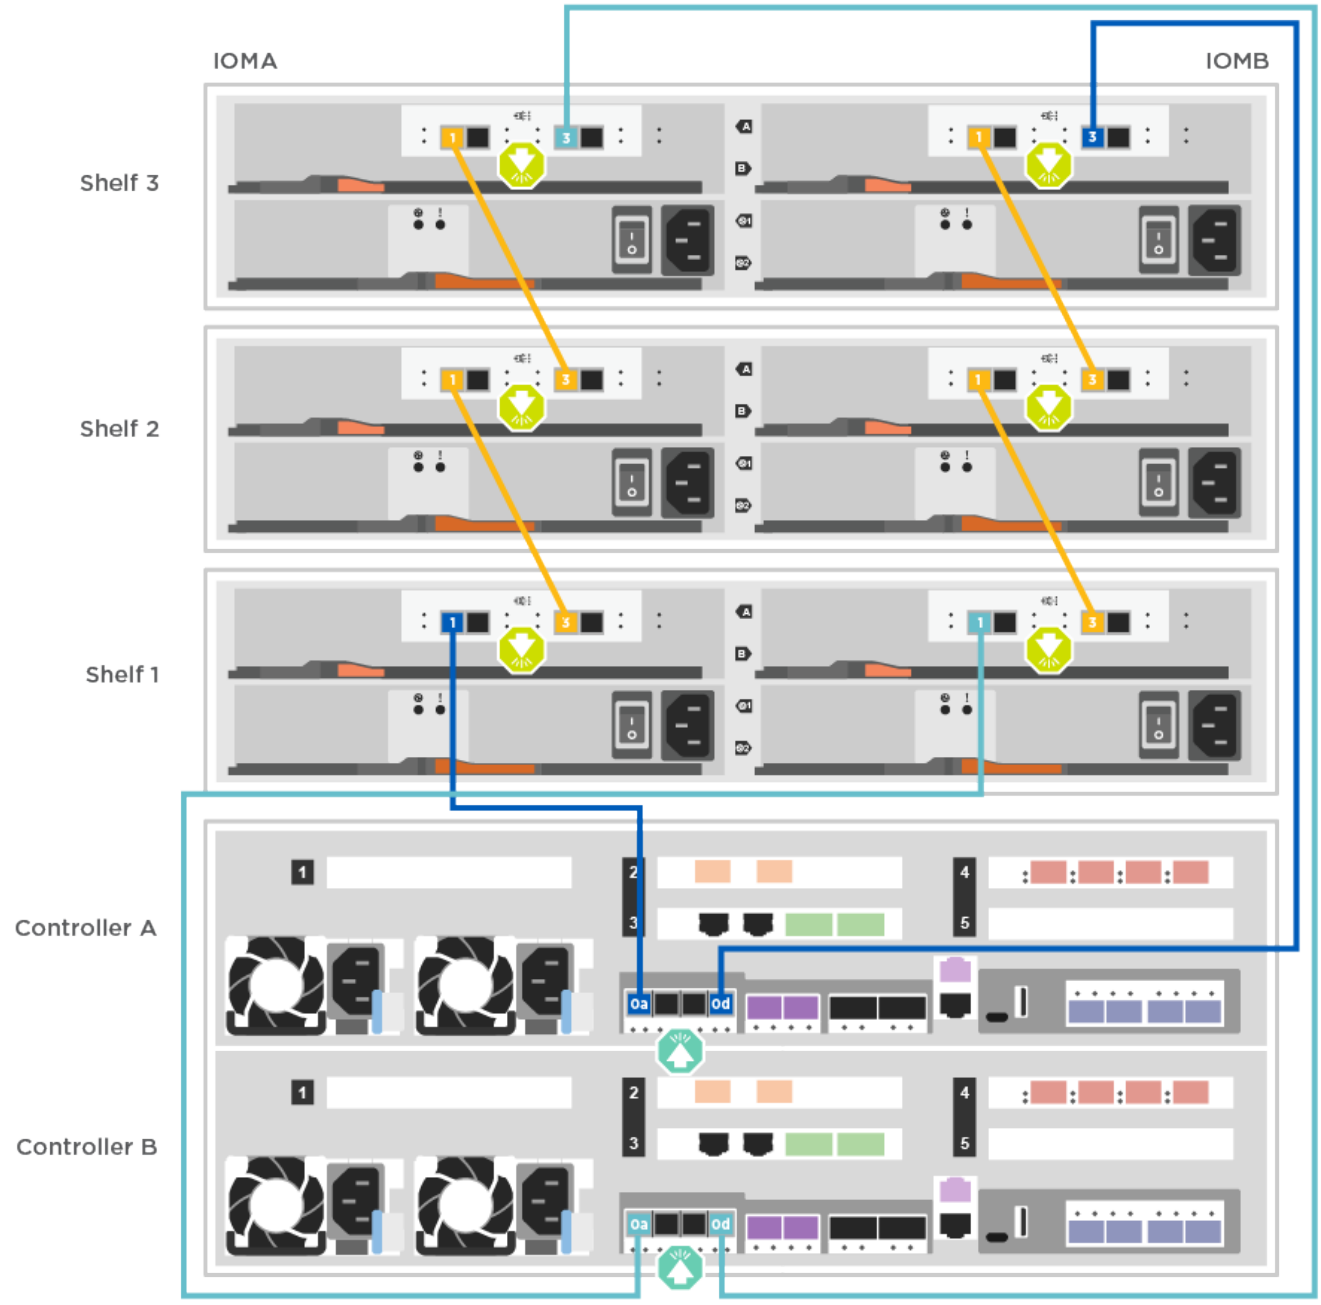

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Completing system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

1. Use the following animation to set one or more drive shelf IDs:

   If your system has NS224 drive shelves, the shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located.

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

5. Select an ONTAP icon listed to discover:

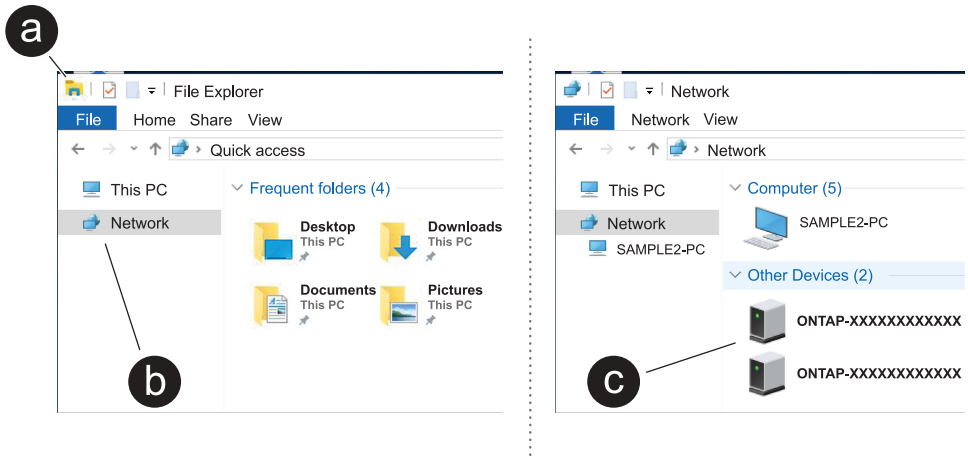

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

   XXXXX is the system serial number for the target node.

   System Manager opens.

6. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

   ONTAP Configuration Guide

7. Set up your account and download Active IQ Config Advisor:

a. Log in to your existing account or create an account.

   NetApp Support Registration

b. Register your system.

   NetApp Product Registration

c. Download Active IQ Config Advisor.

   NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet .
   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Use the following animation to set one or more drive shelf IDs:

   Animation - Set drive shelf IDs

   If your system has NS224 drive shelves, the shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located.

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   FAS8300 and FAS8700 shown.

   Animation - Power on the controllers

   Initial booting may take up to eight minutes.

4. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the *NetApp ONTAP Configuration guide*.

      ONTAP Configuration Guide

6. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain ASA A400 hardware

Maintain the hardware of your ASA A400 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the ASA A400 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the ASA A400 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.
NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.
PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - ASA A400

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your ASA A400 system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - ASA A400

Before replacing the boot media in your ASA A400, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery

process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - ASA A400

Shut down the impaired controller in your ASA A400 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - ASA A400

The boot media in your ASA A400 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

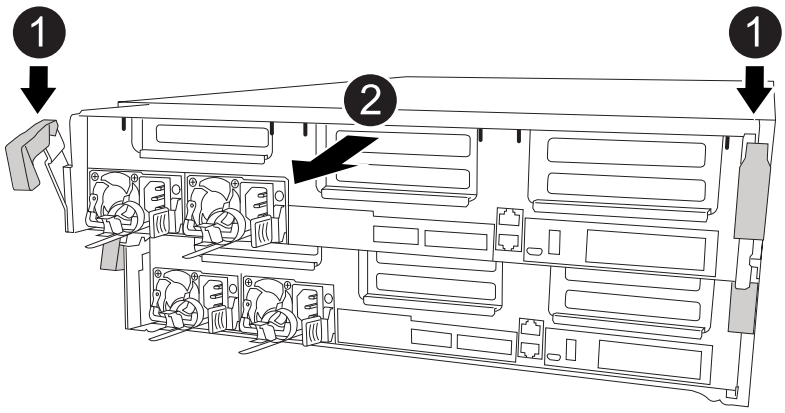

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. Open the air duct:

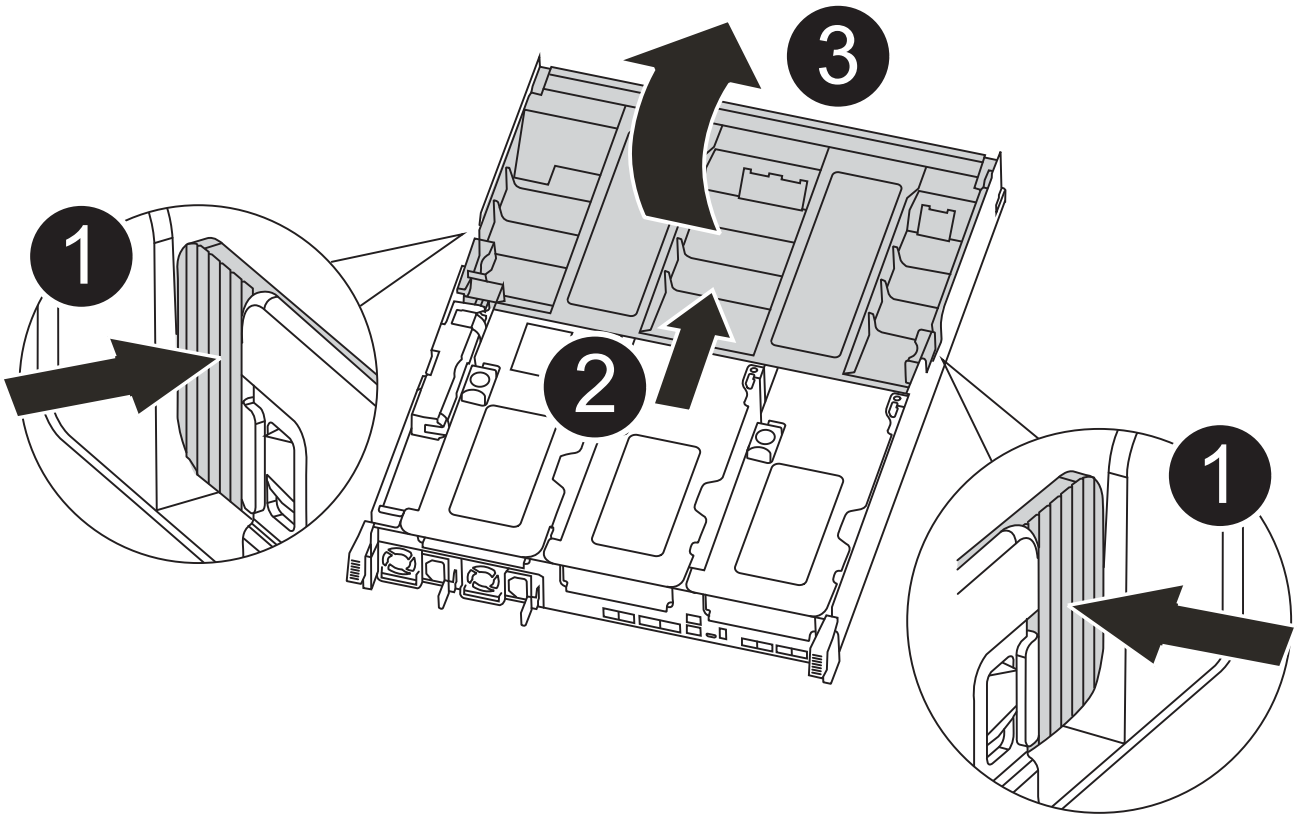

1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.

b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

9. Locate and remove the boot media from the controller module:

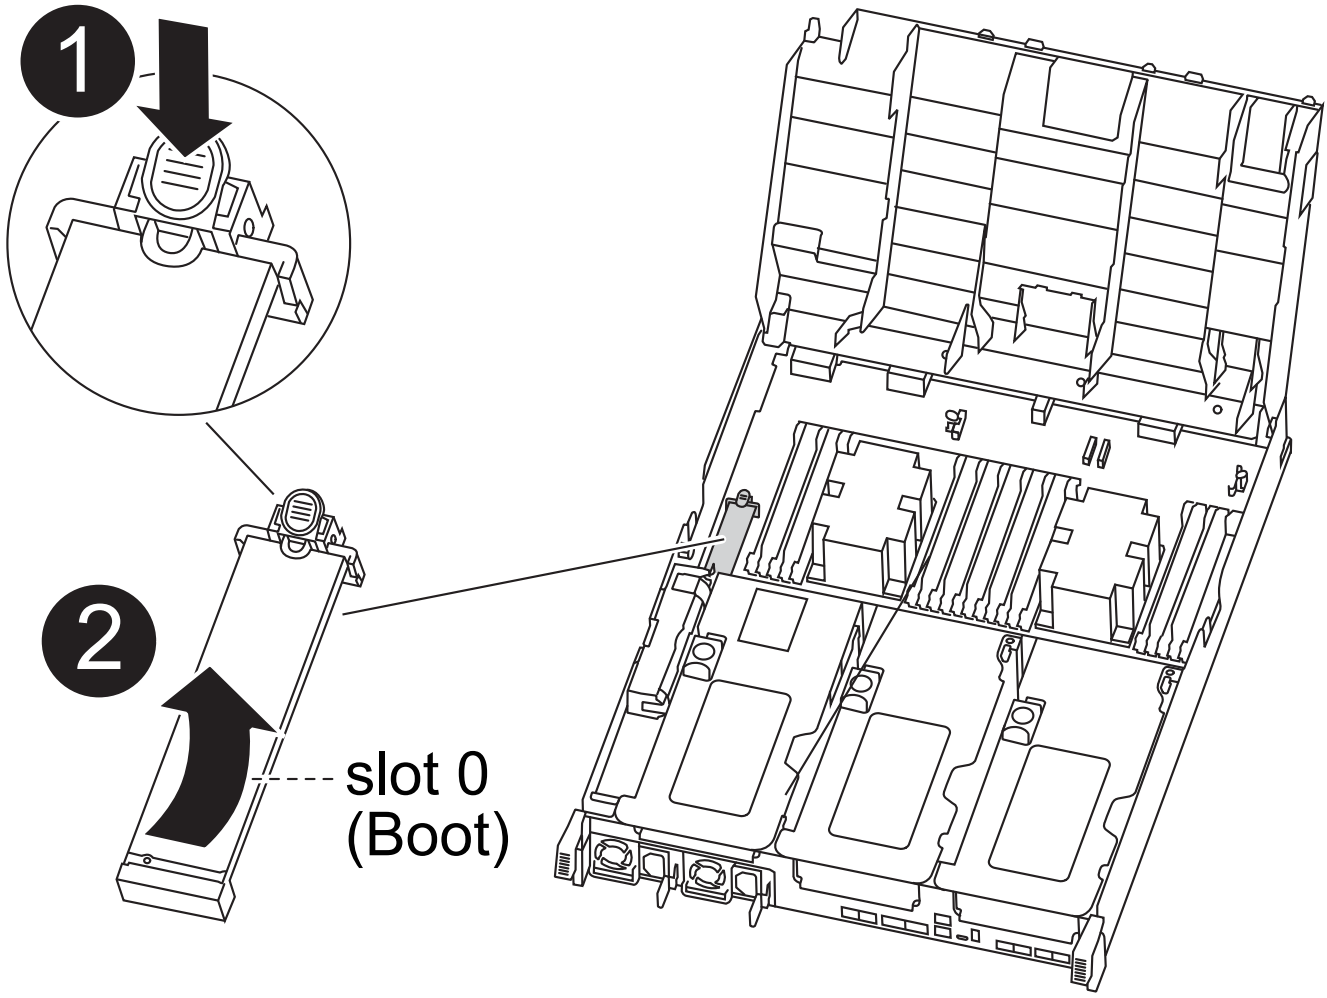

1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.

b. Rotate the boot media up and gently pull the boot media out of the socket.

10. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
11. Check the boot media to make sure that it is seated squarely and completely in the socket.

    If necessary, remove the boot media and reseat it into the socket.

12. Lock the boot media in place:
    a. Rotate the boot media down toward the motherboard.
    b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
    c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
13. Close the air duct.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - ASA A400

After installing the new boot media device in your ASA A400 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - ASA A400

If a component in your ASA A400 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - ASA A400

Get started with replacing the boot media in your ASA A400 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - ASA A400

Before replacing the boot media in your ASA A400 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - ASA A400

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - ASA A400

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster configuration

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 3: Controller is in a two-node Metrocluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes

that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the boot media and prepare for manual boot recovery - ASA A400

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

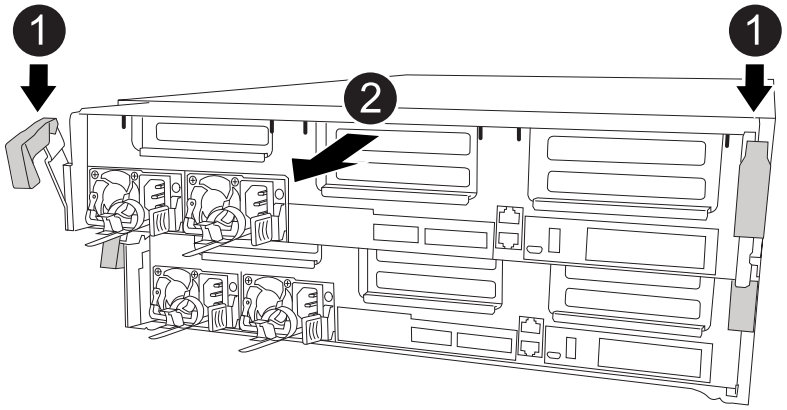

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 2: Replace the boot media

You must locate the boot media in the controller module (see the FRU map on the controller module), and then follow the directions to replace it.

Before you begin

Although the contents of the boot media is encrypted, it is a best practice to erase the contents of the boot media before replacing it. For more information, see the Statement of Volatility for your system on the NetApp Support Site.

> You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the boot media.

Animation - Replace the boot media

Steps

1. Open the air duct:

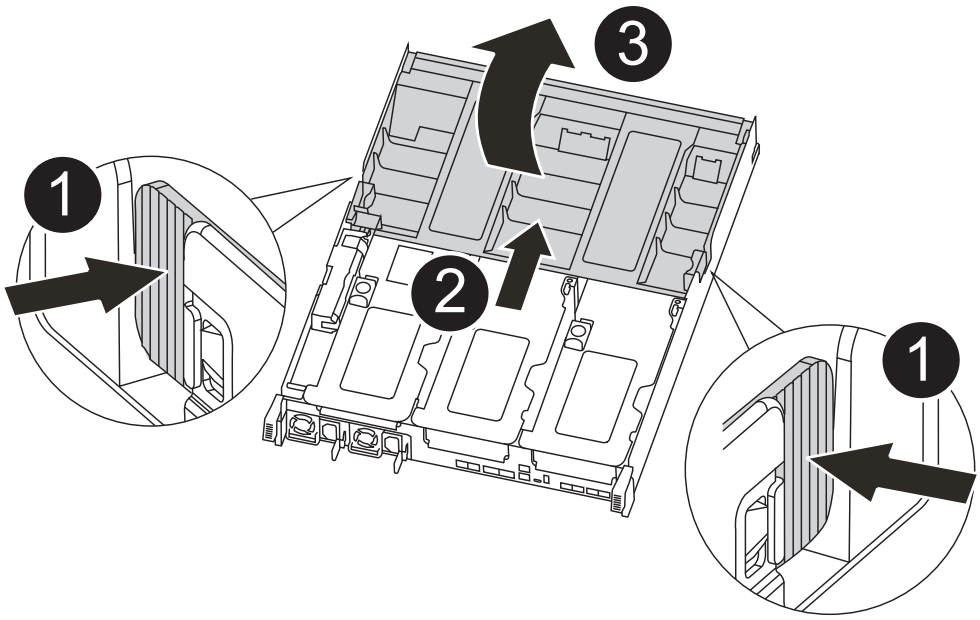

1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

2. Locate and remove the boot media from the controller module:

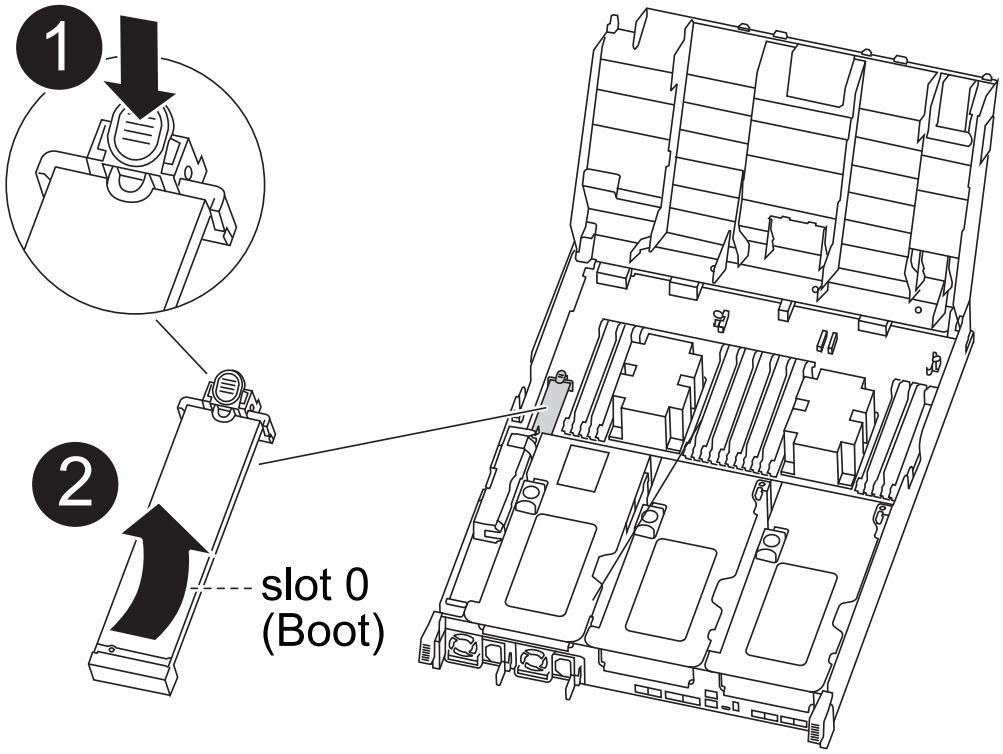

1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
b. Rotate the boot media up and gently pull the boot media out of the socket.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Lock the boot media in place:
   a. Rotate the boot media down toward the motherboard.
   b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
   c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
6. Close the air duct.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed does not have a boot image, so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the `var` file system.

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      There are two folders in the unzipped service image file:

      - `boot`
      - `efi`

   c. Copy the `efi` folder to the top directory on the USB flash drive.

      If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

      The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

   d. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.
5. Plug the power cable into the power supply and reinstall the power cable retainer.
6. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.
7. Complete the installation of the controller module:

a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

   The locking latches rise when the controller module is fully seated.

   > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

d. If you have not already done so, reinstall the cable management device.

8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then `halt` the controller to boot to LOADER.

9. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:

   a. Boot to Maintenance mode: `boot_ontap maint`

   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t` *`initiator adapter_name`*

   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - ASA A400

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.

3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `...would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore OKM, NSE, and NVE - ASA A400

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - ASA A400

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - ASA A400

To replace the chassis, you must move the fans and controller modules from the impaired chassis to the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial

outage in a multinode cluster.

Shut down the controllers - ASA A400

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Shut down the controllers when replacing a chassis

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Shut down a controller in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the

`-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Move the fans, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move the fans

To move the fan modules to the replacement chassis when replacing the chassis, you must perform a specific sequence of tasks.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.
4. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.
5. Set the fan module aside.
6. Repeat the preceding steps for any remaining fan modules.

7. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.
10. Repeat these steps for the remaining fan modules.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
   b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete chassis replacement - ASA A400

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for *HA-state* can be one of the following:

      - `ha`
      - `mcc`
      - `mcc-2n`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

4. Reinstall the bezel on the front of the system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller replacement - ASA A400

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement node* is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - ASA A400

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the controller - ASA A400

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

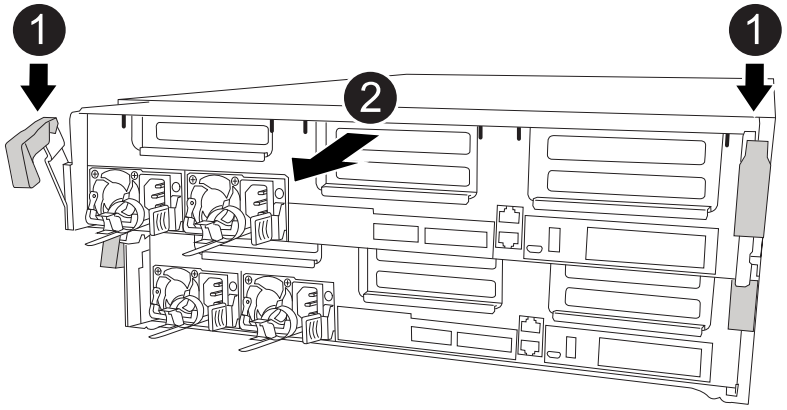

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. On the replacement controller module, open the air duct and remove the empty risers from the controller module using the animation, illustration, or the written steps:

Animation - Remove the empty risers from the replacement controller module

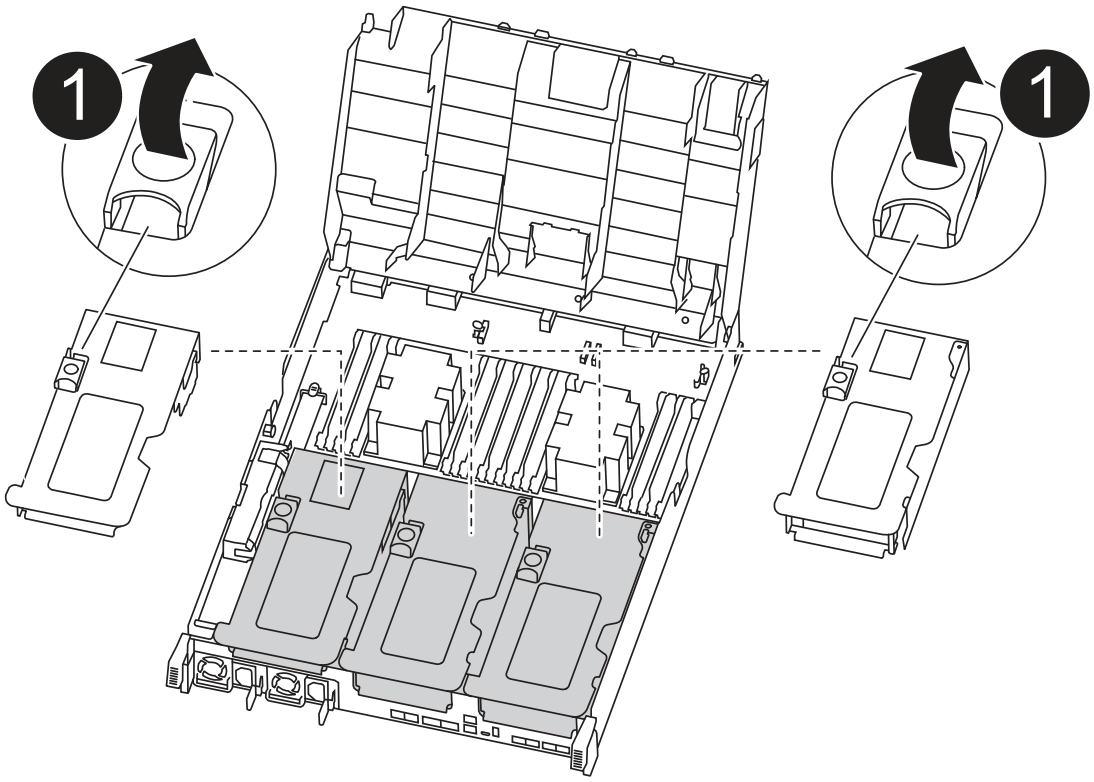

1	Riser latches

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
c. Rotate the riser locking latch on the left side of riser 1 up and toward air duct, lift the riser up, and then set it aside.
d. Repeat the previous step for the remaining risers.

Step 2: Move the power supplies

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

You can use the following animation, illustration, or the written steps to move the power supplies to the replacement controller module.

Animation - Move the power supplies

1. Remove the power supply:

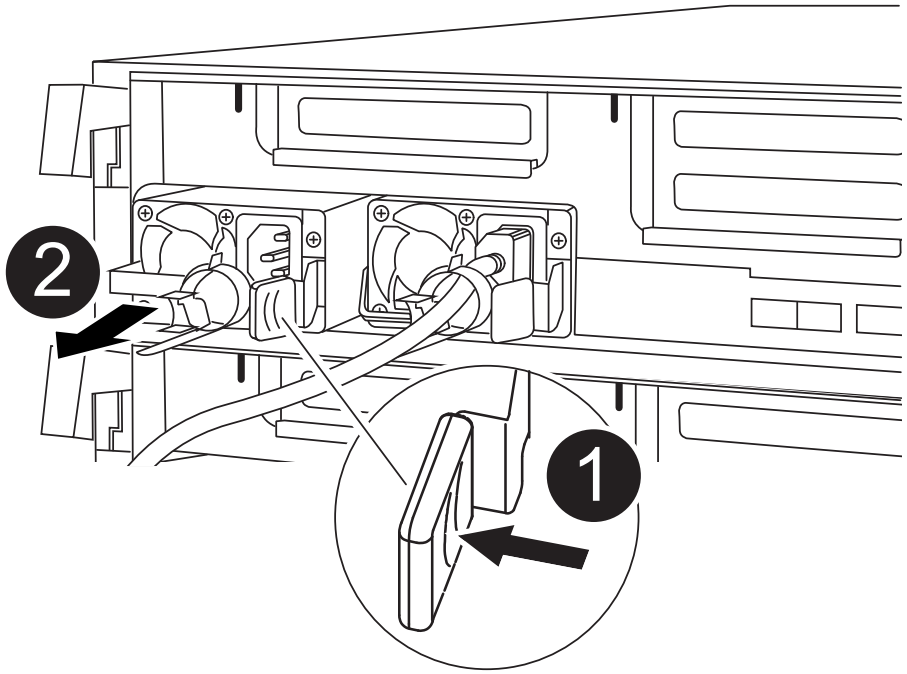

1	PSU locking tab
2	Power cable retainer

a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.

b. Press the blue locking tab to release the power supply from the chassis.

c. Using both hands, pull the power supply out of the chassis, and then set it aside.

1. Move the power supply to the new controller module, and then install it.
2. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

3. Repeat the preceding steps for any remaining power supplies.

Step 3: Move the NVDIMM battery

To move the NVDIMM battery from the impaired controller module to the replacement controller module, you must perform a specific sequence of steps.

You can use the following animation, illustration, or the written steps to move the NVDIMM battery from the impaired controller module to the replacement controller module.

Animation - Move the NVDIMM battery

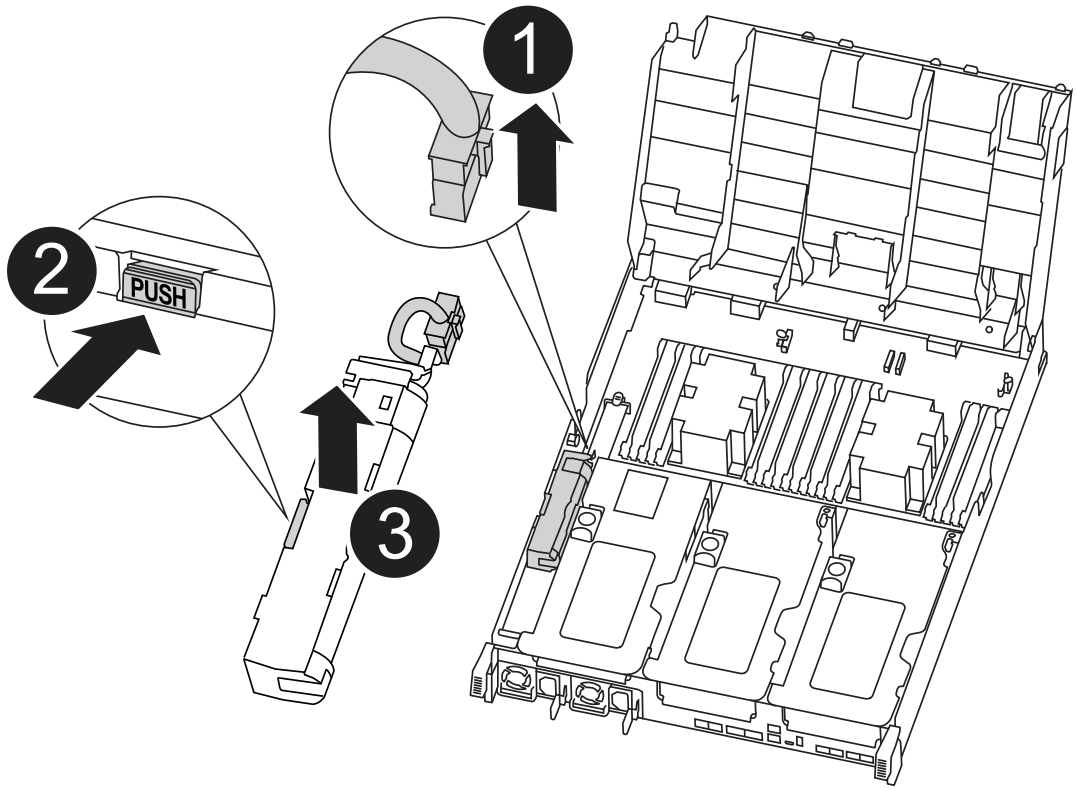

1	NVDIMM battery plug
2	NVDIMM battery locking tab
3	NVDIMM battery

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.
3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.

Do not plug the battery cable back into the motherboard until instructed to do so.

Step 4: Move the boot media

You must locate the boot media, and then follow the directions to remove it from the impaired controller module and insert it into the replacement controller module.

You can use the following animation, illustration, or the written steps to move the boot media from the impaired controller module to the replacement controller module.

Animation - Move the boot media

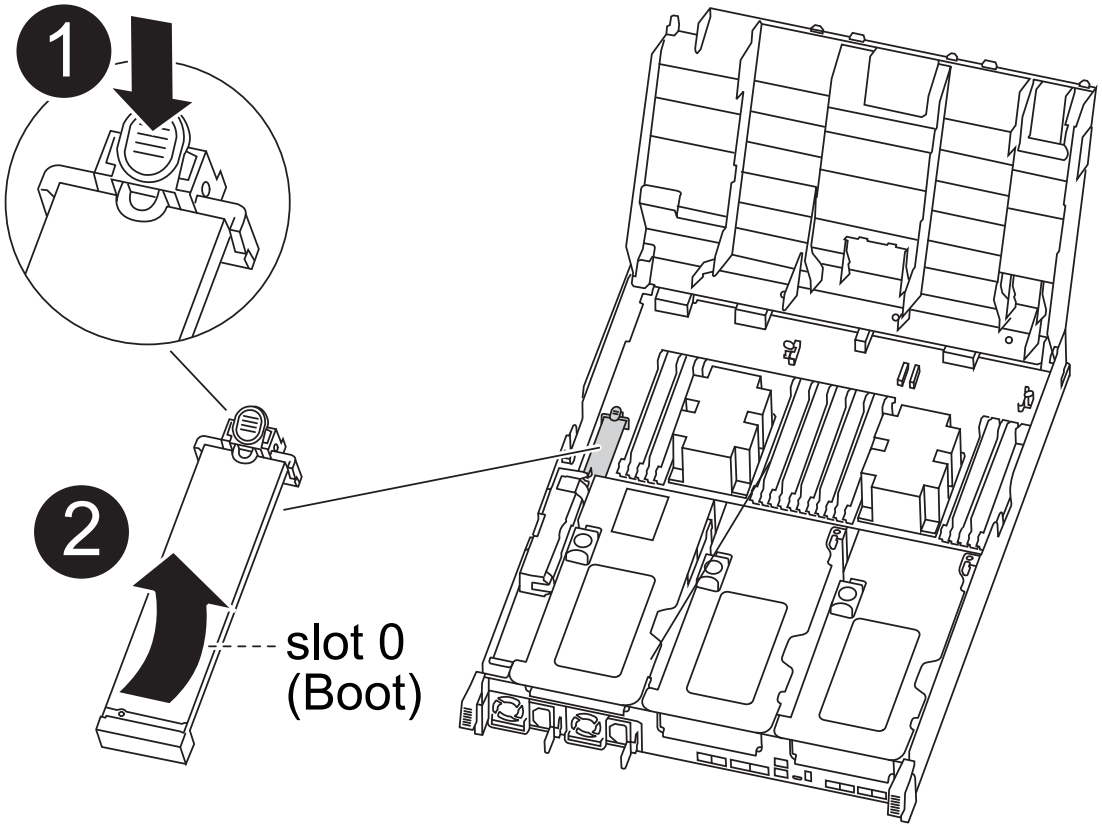

1	Boot media locking tab
2	Boot media

1. Locate and remove the boot media from the controller module:
   a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
   b. Rotate the boot media up and gently pull the boot media out of the socket.
2. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
3. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

4. Lock the boot media in place:

a. Rotate the boot media down toward the motherboard.

b. Press the blue locking button so that it is in the open position.

c. Placing your fingers at the end of the boot media by the blue button, firmly push down on the boot media end to engage the blue locking button.

Step 5: Move the PCIe risers and mezzanine card

As part of the controller replacement process, you must move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

You can use the following animations, illustrations, or the written steps to move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

Moving PCIe riser 1 and 2 (left and middle risers):

Animation - Move PCI risers 1 and 2

Moving the mezzanine card and riser 3 (right riser):

Animation - Move the mezzanine card and riser 3

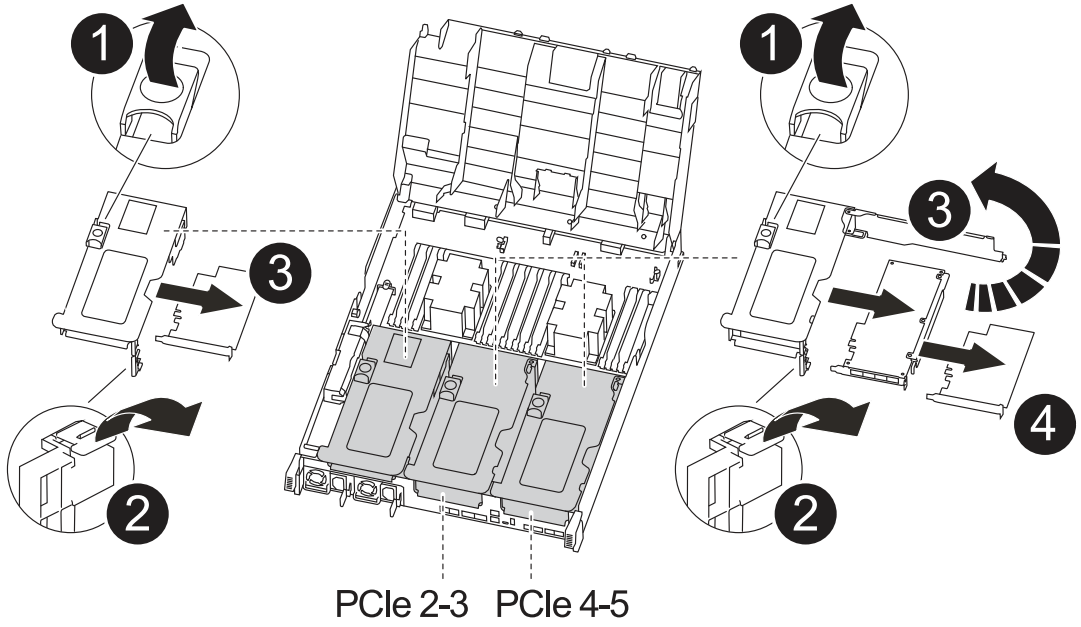

1	Riser locking latch
2	PCI card locking latch
3	PCI locking plate
4	PCI card

1. Move PCIe risers one and two from the impaired controller module to the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then move it to the replacement controller module.
   d. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins, push the riser squarely into the socket on the motherboard, and then rotate the latch down flush with the sheet metal on the riser.
   e. Repeat this step for riser number 2.
2. Remove riser number 3, remove the mezzanine card, and install both into the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then set it aside on a stable, flat surface.
   d. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket, and then move it to the replacement controller module.
   e. Install the mezzanine in the replacement controller and secure it with the thumbscrews.
   f. Install the third riser in the replacement controller module.

Step 6: Move the DIMMs

You need to locate the DIMMs, and then move them from the impaired controller module to the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

You can use the following animation, illustration, or the written steps to move the DIMMs from the impaired controller module to the replacement controller module.

Animation - Move the DIMMs

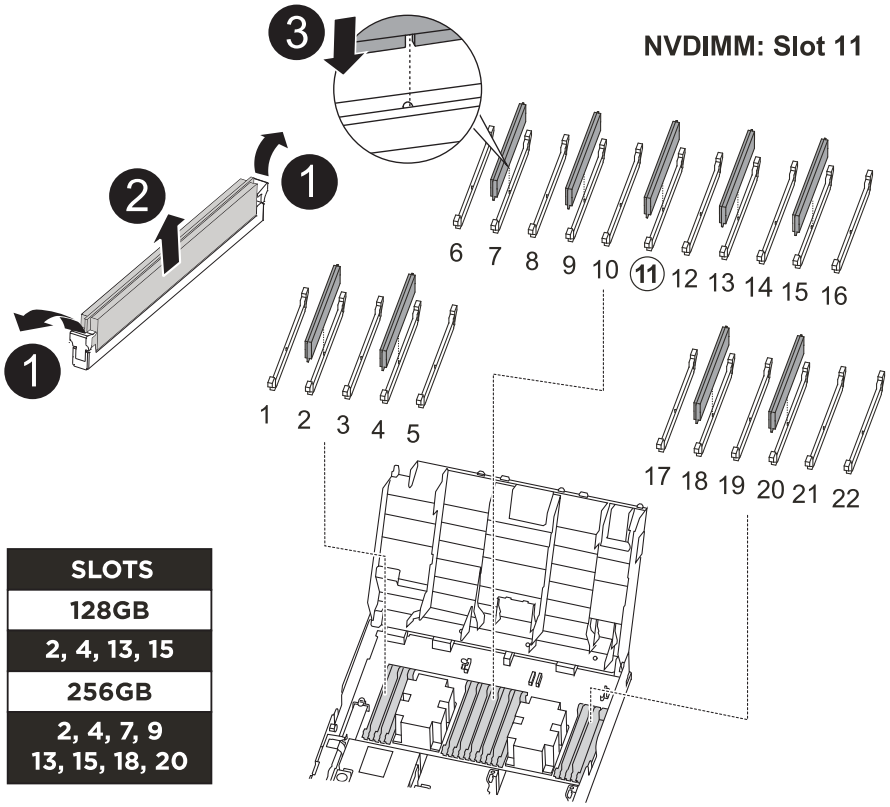

1	DIMM locking tabs
2	DIMM
3	DIMM socket

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Verify that the NVDIMM battery is not plugged into the new controller module.
4. Move the DIMMs from the impaired controller module to the replacement controller module:

   Make sure that you install the each DIMM into the same slot it occupied in the impaired controller module.

   a. Eject the DIMM from its slot by slowly pushing apart the DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

      Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   b. Locate the corresponding DIMM slot on the replacement controller module.

   c. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the

DIMM squarely into the socket.

The DIMMs fit tightly in the socket, but should go in easily. If not, realign the DIMM with the socket and reinsert it.

d. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

e. Repeat these substeps for the remaining DIMMs.

5. Plug the NVDIMM battery into the motherboard.

Make sure that the plug locks down onto the controller module.

Step 7: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

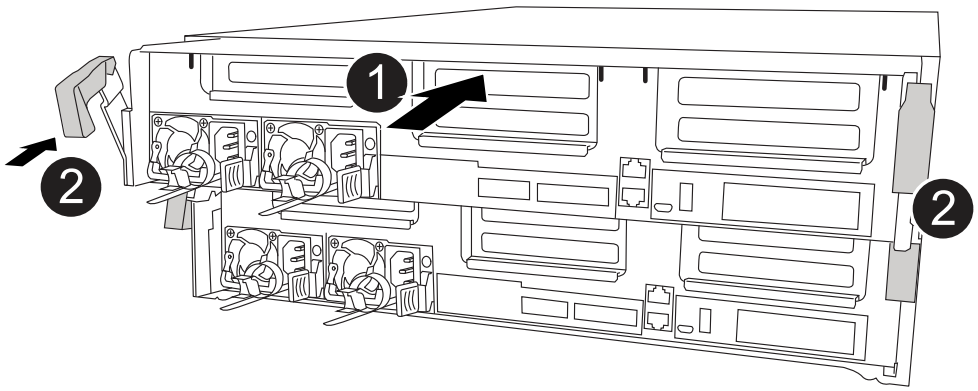

1	Slide controller into the chassis
2	Locking latches

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:

a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

Restore and verify the system configuration - ASA A400

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - ASA A400

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.

c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.

d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`

   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *`local-node-name`* `partner savecore -s`

d. Return to the admin privilege level: `set -privilege admin`

5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at

which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The *replacement* controller is the current owner of the disks on the disaster site.

     Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - ASA A400

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no

configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    ------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace DIMM - ASA A400

You must replace a DIMM in the controller module when your system registers an increasing number of correctable error correction codes (ECC); failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

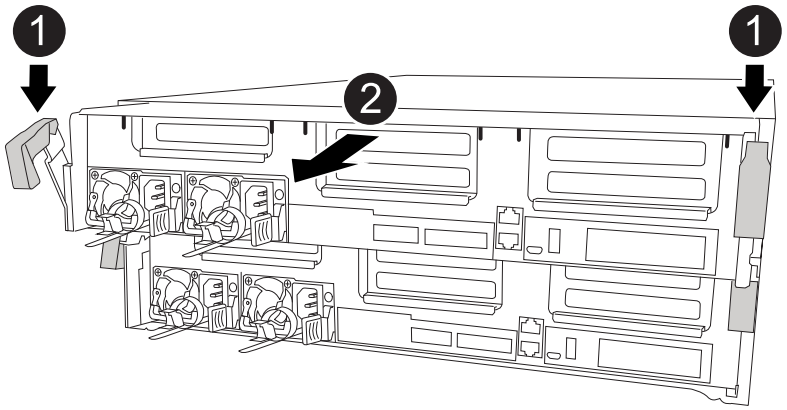

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace system DIMMs

Replacing a system DIMM involves identifying the target DIMM through the associated error message, locating the target DIMM using the FRU map on the air duct, and then replacing the DIMM.

You can use the following animation, illustration, or the written steps to replace a system DIMM.

> The animation and illustration show empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace a system DIMM

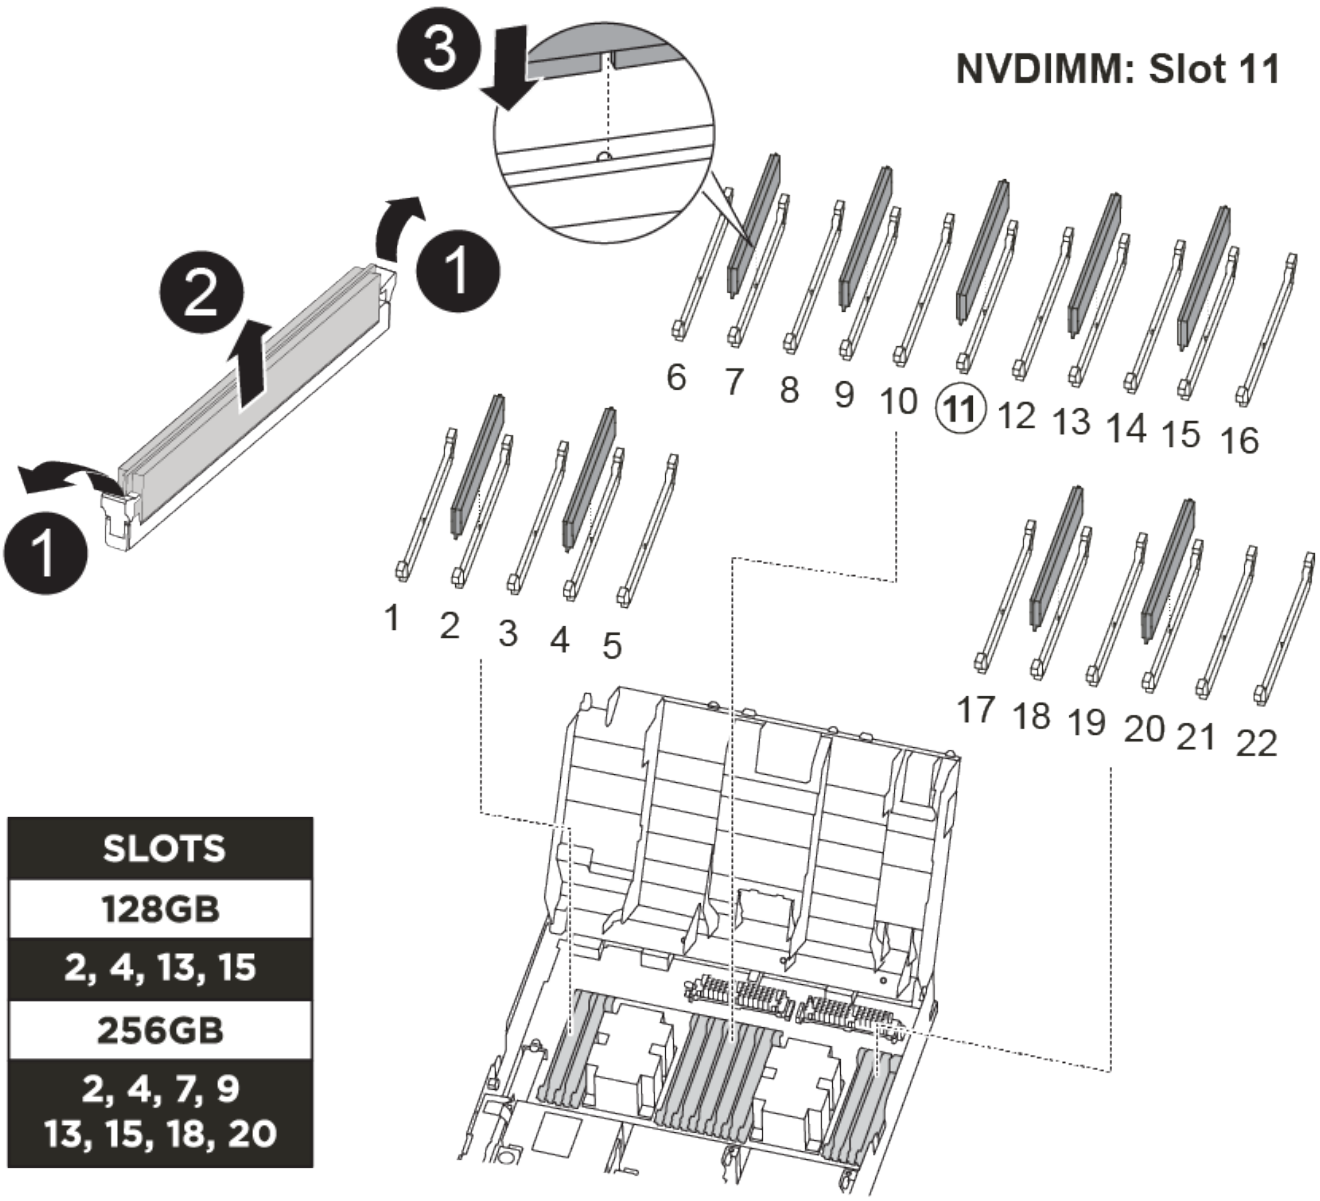

1	DIMM locking tabs
2	DIMM
3	DIMM socket

The DIMMs are located in sockets 2, 4, 13, and 15. The NVDIMM is located in slot 11.

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely

open position.

2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
4. Eject the DIMM from its socket by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the socket.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

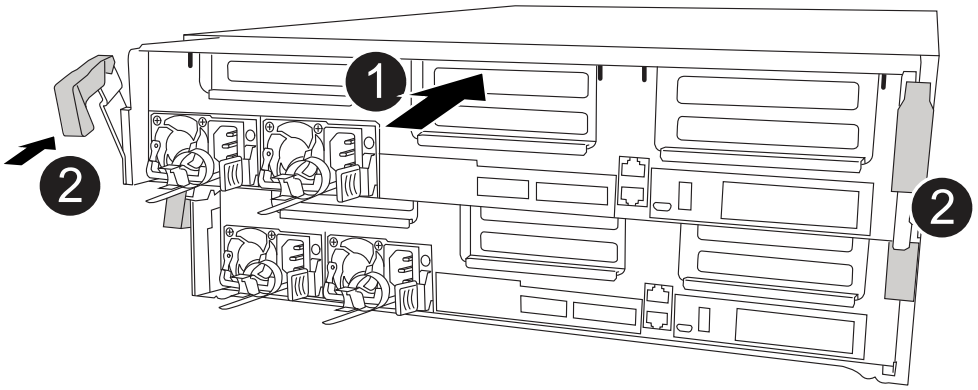

1	Controller module
2	Controller locking latches

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
   b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.
   g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto`

```
-giveback true
```

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    ------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a fan module - ASA A400

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

You can use the following animation, illustration, or the written steps to hot-swap a fan module.

Animation - Replace a fan

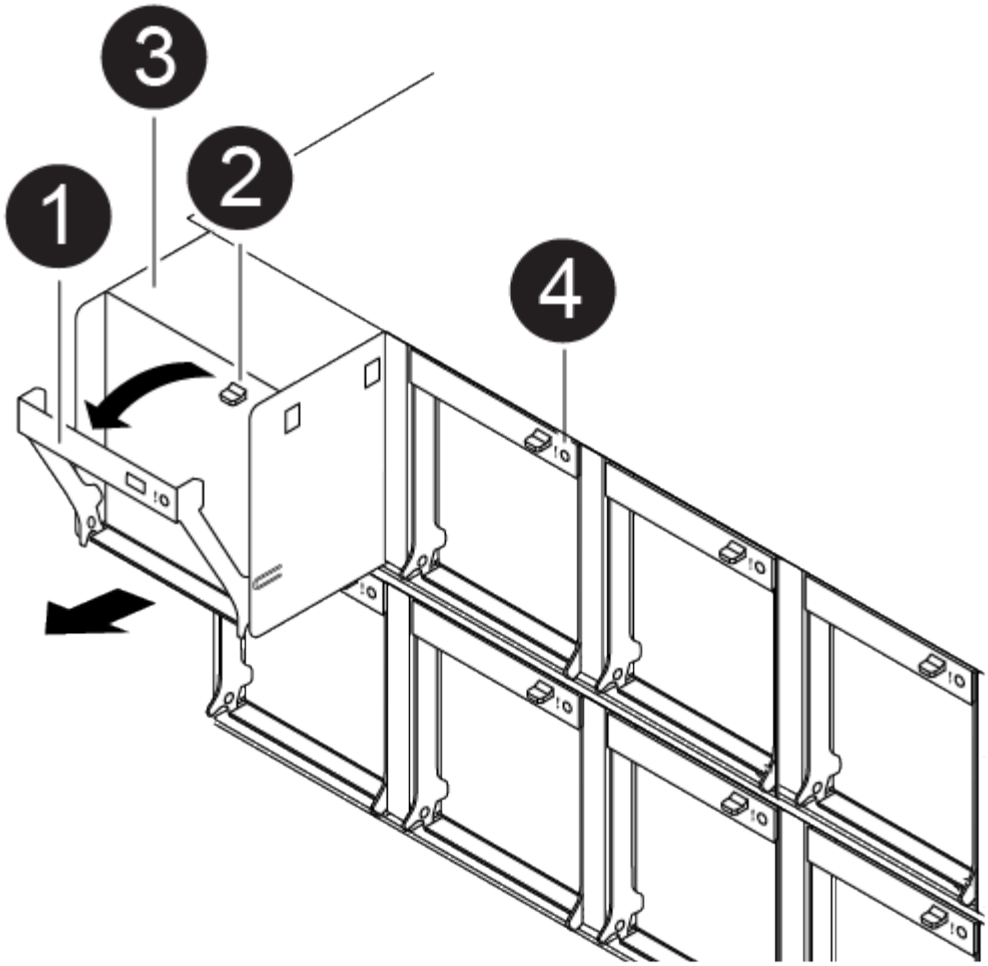

1	Fan handle
2	Locking tab
3	Fan
4	Status LED

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.

5. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

6. Set the fan module aside.
7. Insert the replacement fan module into the chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The Attention LED should not be lit after the fan is seated and has spun up to operational speed.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - ASA A400

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

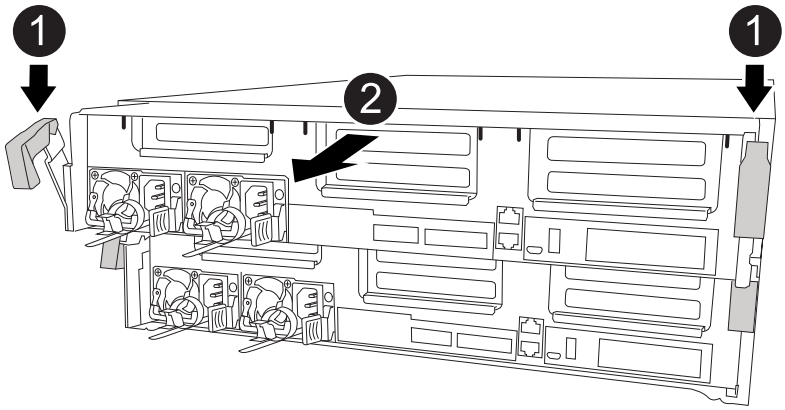

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module. See the FRU map inside the controller module to locate the NVDIMM battery.

The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.

You can use the following animation, illustration, or the written steps to replace the NVDIMM battery.

Animation - Replace the NVDIMM battery

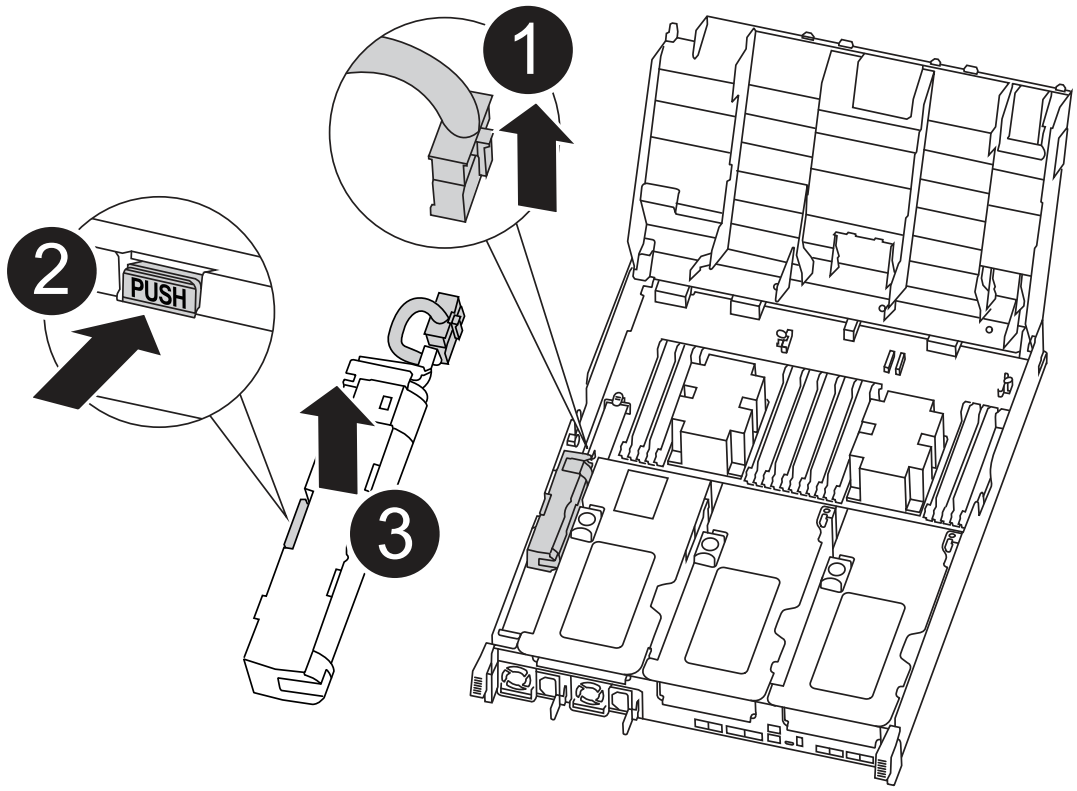

1	Battery plug
2	Locking tab
3	NVDIMM battery

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.
3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Remove the replacement battery from its package.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.
7. Plug the battery plug back into the controller module, and then close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

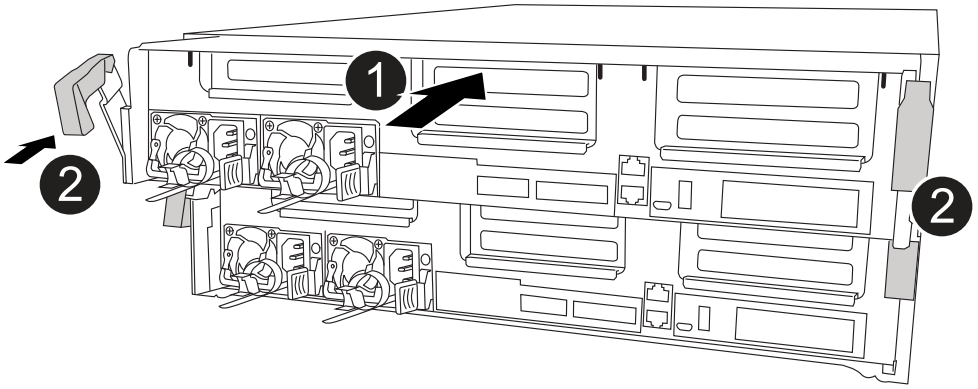

1	Controller module
2	Controller locking latches

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
   b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to

interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace NVDIMM - ASA A400

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

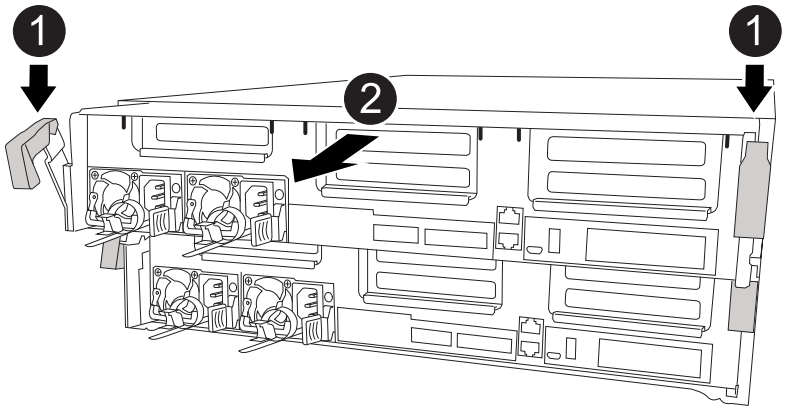

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the FRU map on top of the air duct or the FRU Map on the top of the slot 1 riser.

- The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.
- Although the contents of the NVDIMM is encrypted, it is a best practice to erase the contents of the NVDIMM before replacing it. For more information, see the Statement of Volatility on the NetApp Support

Site.

You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the NVDIMM.

The animation shows empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace the NVDIMM

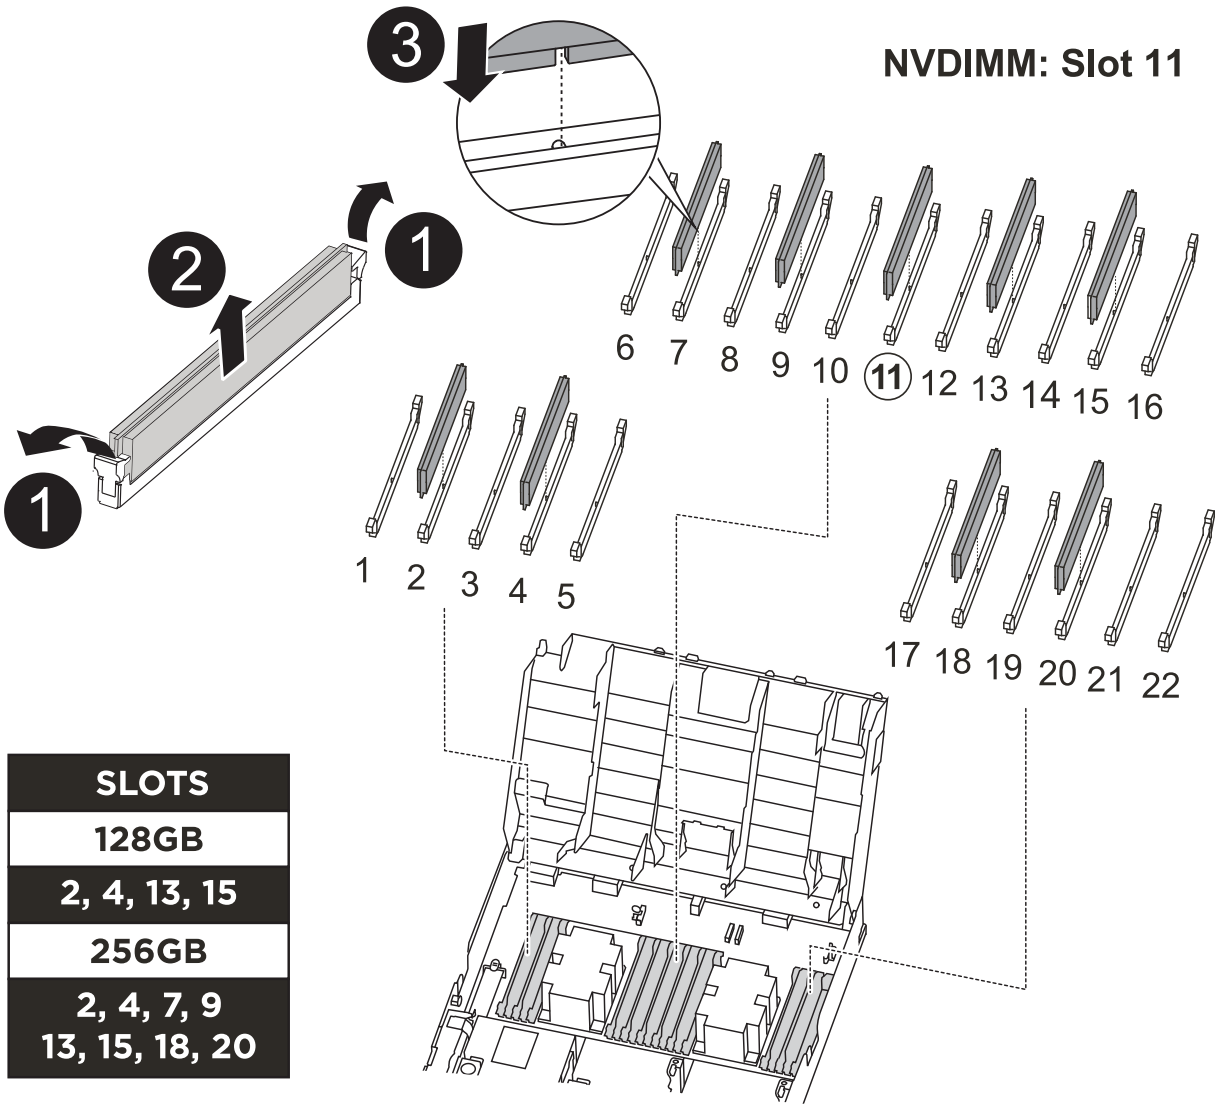

1	DIMM locking tabs
2	DIMM

3	DIMM socket

1. Open the air duct and then locate the NVDIMM in slot 11 on your controller module.

   The NVDIMM looks significantly different than system DIMMs.

2. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

3. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

4. Locate the slot where you are installing the NVDIMM.
5. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
7. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe or mezzanine card - ASA A400

To replace a PCIe or mezzanine card, you must disconnect the cables and any SFP and QSFP modules from the cards, replace the failed PCIe or mezzanine card, and then recable the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft

vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

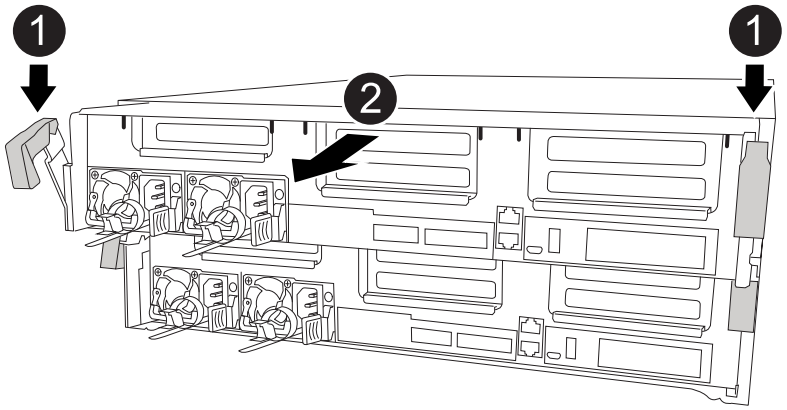

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace a PCIe card

To replace a PCIe card, you must locate the failed PCIe card, remove the riser that contains the card from the controller module, replace the card, and then reinstall the PCIe riser in the controller module.

You can use the following animation, illustration, or the written steps to replace a PCIe card.

Animation - Replace a PCIe card

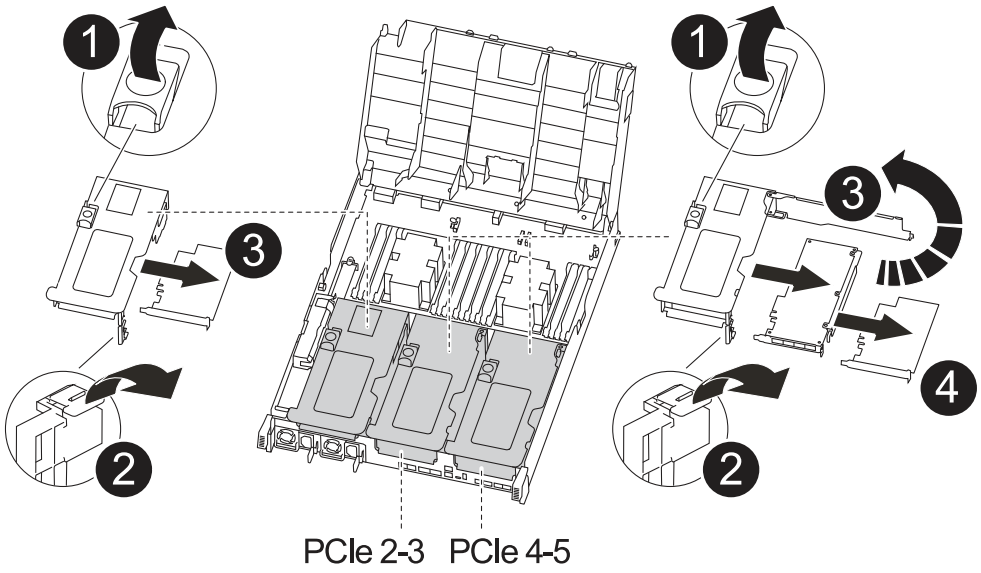

1	Riser locking latch
2	PCI card locking latch
3	PCI locking plate
4	PCI card

1. Remove the riser containing the card to be replaced:
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up straight up and set it aside on a stable flat surface,
2. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. For risers 2 and 3 only, swing the side panel up.
   d. Remove the PCIe card from the riser by gently pushing up on the bracket and lift the card straight out of the socket.
3. Install the replacement PCIe card in the riser by aligning the card with the socket, press the card into the socket and then close the side panel on the riser, if present.

Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The PCIe card must be fully and evenly seated in the slot.

If you are installing a card in the bottom slot and cannot see the card socket well, remove the top card so that you can see the card socket, install the card, and then reinstall the card you removed from the top slot.

4. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 4: Replace the mezzanine card

The mezzanine card is located under riser number 3 (slots 4 and 5). You must remove that riser to access the mezzanine card, replace the mezzanine card, and then reinstall riser number 3. See the FRU map on the controller module for more information.

You can use the following animation, illustration, or the written steps to replace the mezzanine card.

Animation - Replace the mezzanine card

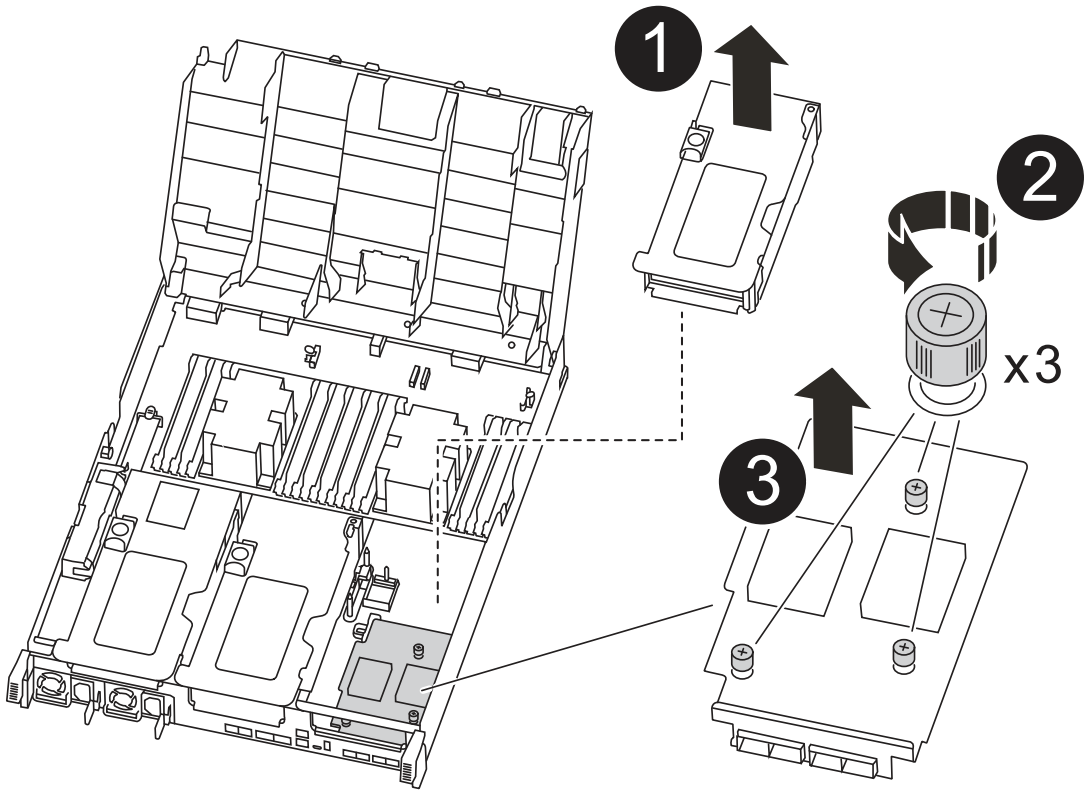

1	PCI riser

2	Riser thumbscrew
3	Riser card

1. Remove riser number 3 (slots 4 and 5):
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up, and then set it aside on a stable, flat surface.
2. Replace the mezzanine card:
   a. Remove any QSFP or SFP modules from the card.
   b. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket and set it aside.
   c. Align the replacement mezzanine card over the socket and the guide pins and gently push the card into the socket.
   d. Tighten the thumbscrews on the mezzanine card.
3. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 5: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

   The locking latches rise when the controller module is fully seated.

   > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

   > If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

5. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
6. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Restore the controller module to operation

To restore the controller, you must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 7: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 8: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replacing a power supply - ASA A400

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting the replacement PSU to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

You can use the following illustration with the written steps to replace the power supply.

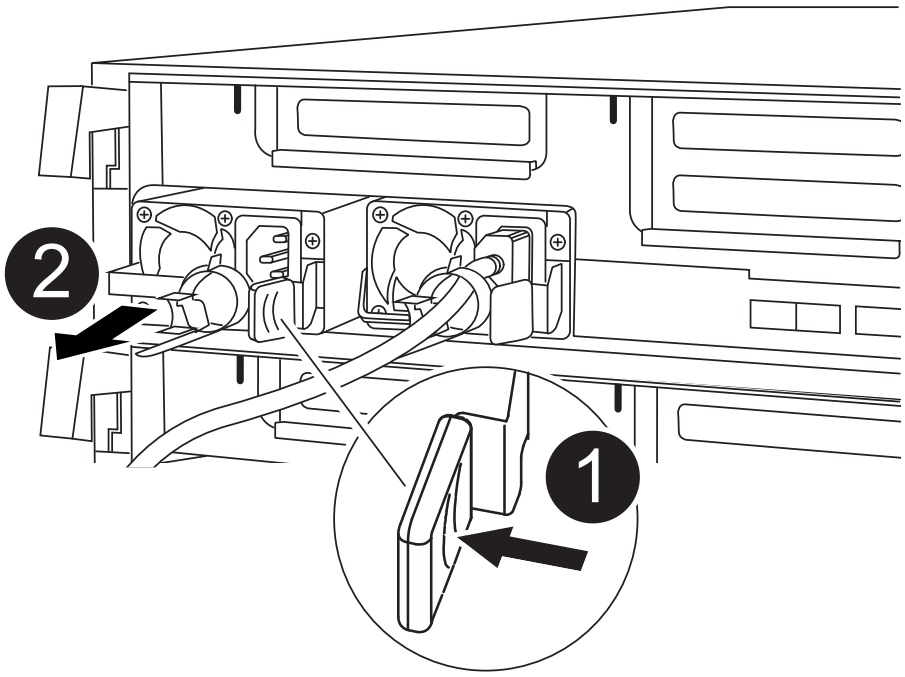

1	PSU locking tab
2	Power cable retainer

1. If you are not already grounded, properly ground yourself.

2. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
3. Disconnect the power supply:
   a. Open the power cable retainer, and then unplug the power cable from the power supply.
   b. Unplug the power cable from the power source.
4. Remove the power supply:
   a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
   b. Press the blue locking tab to release the power supply from the chassis.
   c. Using both hands, pull the power supply out of the chassis, and then set it aside.
5. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

6. Rotate the cam handle so that it is flush against the power supply.
7. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA A400

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

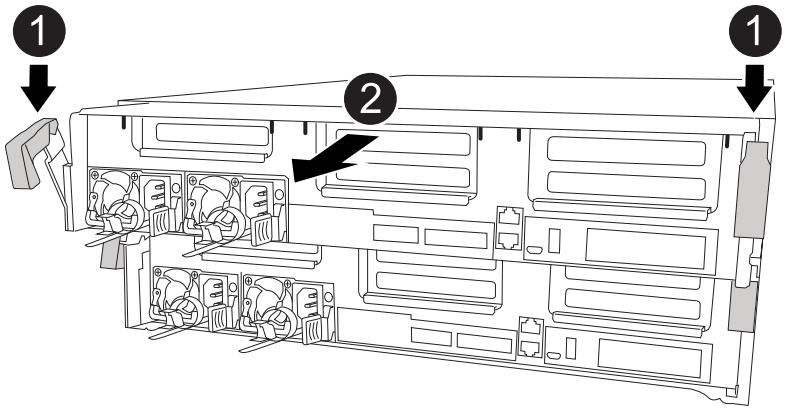

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the RTC battery

You need to locate the RTC battery inside the controller module, and then follow the specific sequence of steps. See the FRU map inside the controller module for the location of the RTC battery.

You can use the following animation, illustration, or the written steps to replace the RTC battery.

Animation- Replace the RTC battery

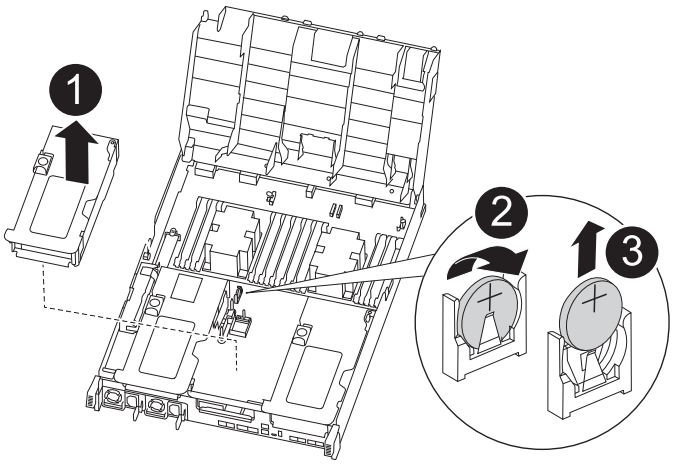

1	Middle riser
2	Remove RTC battery
3	Seat RTC battery

1. If you are not already grounded, properly ground yourself.
2. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
3. Locate, remove, and then replace the RTC battery:
   a. Using the FRU map, locate the RTC battery on the controller module.
   b. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

   > Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

   c. Remove the replacement battery from the antistatic shipping bag.
   d. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
5. Close the air duct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. If you have not already done so, reinstall the cable management device.
   d. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      > If your system stops at the boot menu, select the option to boot to LOADER.

6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA A800 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations (including ASA configurations), you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick steps - ASA A800

This page gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use the AFF A800 Installation and Setup Instructions if you are familiar with installing NetApp systems.

The ASA A800 uses the same installation procedure as the AFF A800 system.

Video steps - ASA A800

The following video shows how to install and cable your new system.

"Animation - Installation and Setup of an AFF A800

The ASA A800 uses the same installation procedure as the AFF A800 system.

Detailed steps - ASA A800

This page gives detailed step-by-step instructions for installing an ASA A800 system.

Step 1: Prepare for installation

To install your system, you need to create an account and register the system. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the NetApp Hardware Universe (HWU) for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

What you need

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

Steps

1. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register (NetApp Product Registration) your system.
2. Download and install NetApp Downloads: Config Advisor on your laptop.
3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Connector type	Part number and length	Type of cable…	For…
100 GbE cable	X66211A-05 (112-00595), 0.5m X66211-1 (112-00573), 1m X66211-2 (112-00574), 2m X66211-5 (112-00576), 5m		• HA interconnect • Cluster interconnect network • Storage, Data

Connector type	Part number and length	Type of cable…	For…
10 GbE cable	X6566B-3-R6 (112-00300), 3m; X6566B-5-R6 (112-00301), 5m		• Data
25 GbE cable	X66240A-2 (112-00598), 2m; X66240A-5 (112-00600), 5m		• Data
RJ-45 (order dependent)	Not applicable		• Management
Fibre Channel	X66250-2 (112-00342) 2m; X66250-5 (112-00344) 5m; X66250-15 (112-00346) 15m; X66250-30 (112-00347) 30m		• Network
Micro-USB console cable	Not applicable		• Console connection during software setup
Power cables	Not applicable		Connecting the PSUs to power source

4. Download and complete the Cluster Configuration Worksheet.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.

   Installing SuperRail into a four-post rack

2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

3. Attach cable management devices (as shown).

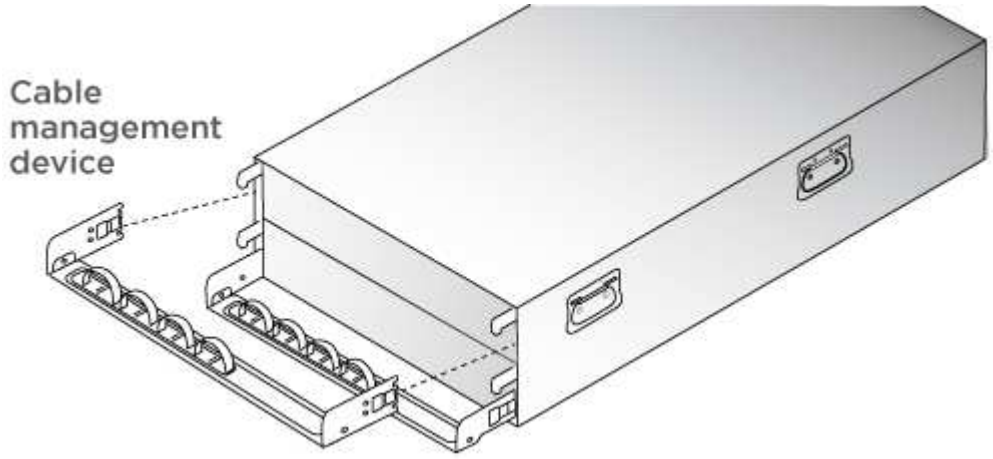

4. Place the bezel on the front of the system.

Step 3: Cable controllers

There is required cabling for your platform's cluster using the two-node switchless cluster method or the cluster interconnect network method. There is optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cable to a host network and storage.

Required cabling: Cable controllers to a cluster

Cable the controllers to a cluster by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Cable a two-node switchless cluster

Management network ports on the controllers are connected to switches. The HA interconnect and cluster interconnect ports are cabled on both controllers.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

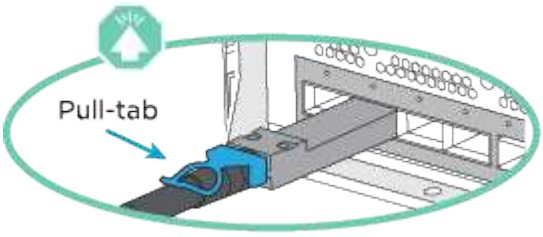

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or the tabulated steps to complete the cabling between the controllers and the switches:

   Animation - Cable a two-node switchless cluster

Step	Perform on each controller module
1	Cable the HA interconnect ports: • e0b to e0b • e1b to e1b 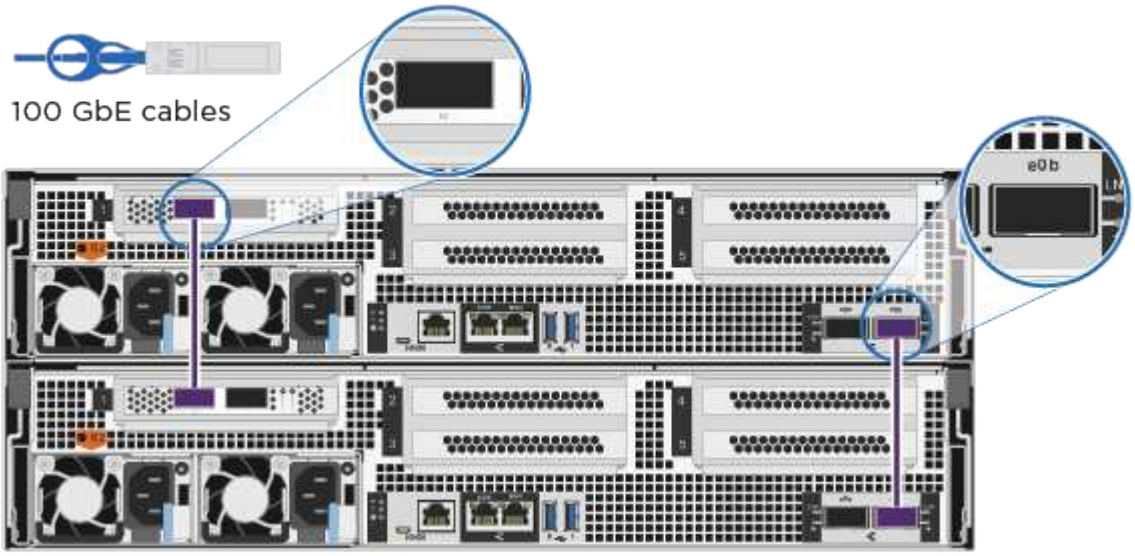
2	Cable the cluster interconnect ports: • e0a to e0a • e1a to e1a 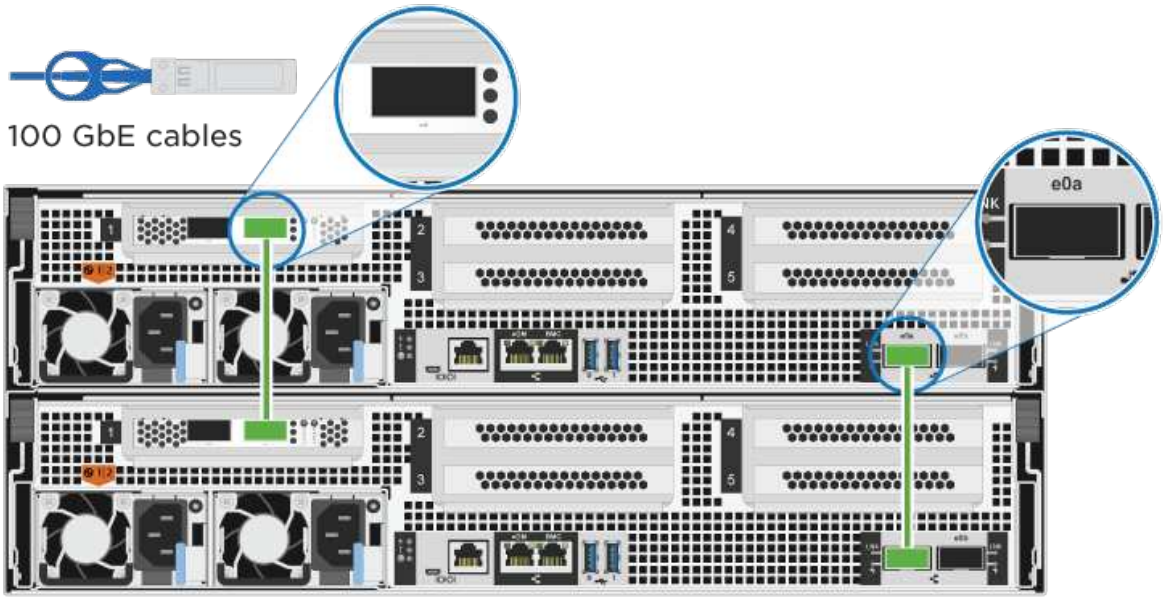

Step	Perform on each controller module
3	Cable the management ports to the management network switches 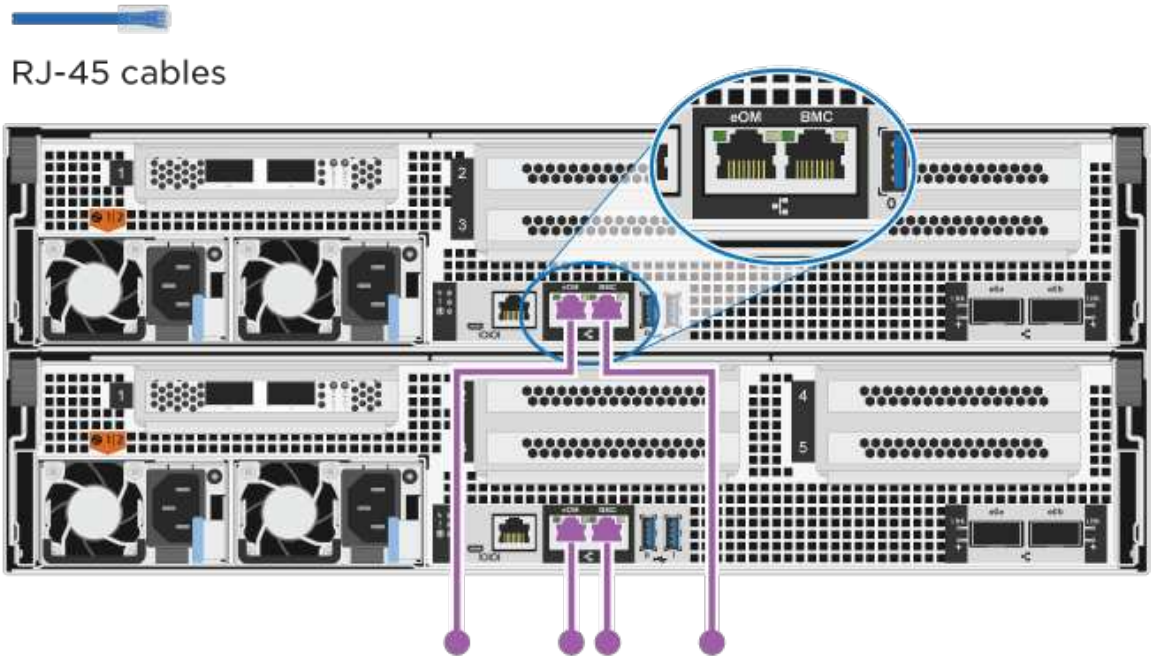
	DO NOT plug in the power cords at this point.

2. To perform optional cabling, see:
   - Option 1: Cable to a Fibre Channel host network
   - Option 2: Cable to a 10GbE host network
   - Option 3: Cable the controllers to a single drive shelf
   - Option 4: Cable the controllers to two drive shelves
3. To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 2: Cable a switched cluster

Cluster interconnect and management network ports on the controllers are connected to switches while the HA interconnect ports are cabled on both controllers.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

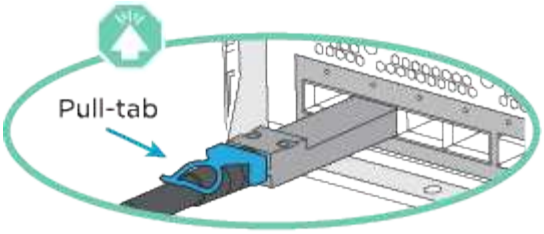

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or the tabulated steps to complete the cabling between the controllers and the switches:

Animation - Cable a switched cluster

Step	Perform on each controller module
1	Cable the HA interconnect ports: • e0b to e0b • e1b to e1b 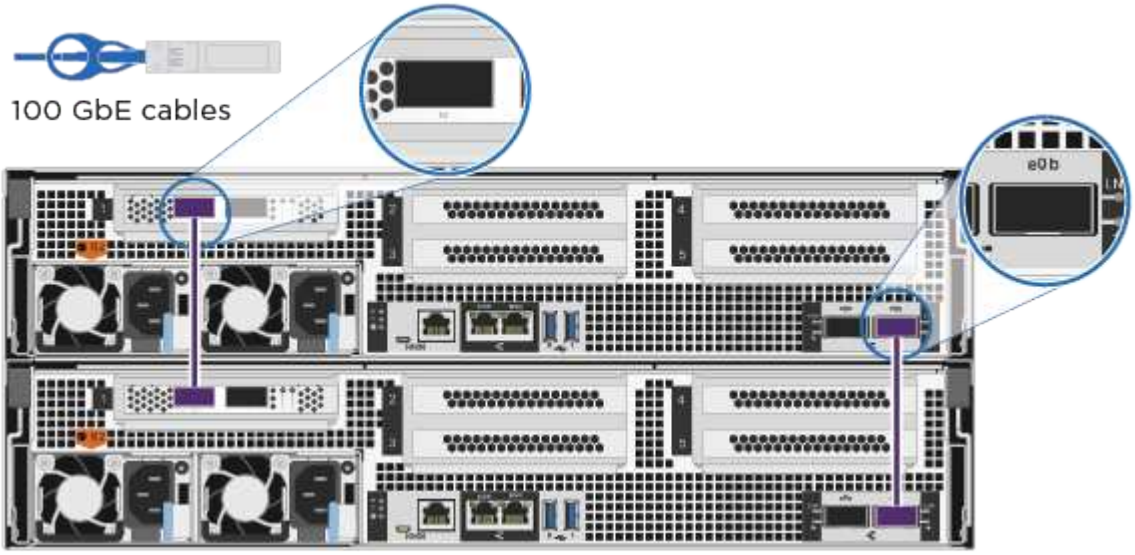

Step	Perform on each controller module
2	Cable the cluster interconnect ports to the 100 GbE cluster interconnect switches. • e0a • e1a 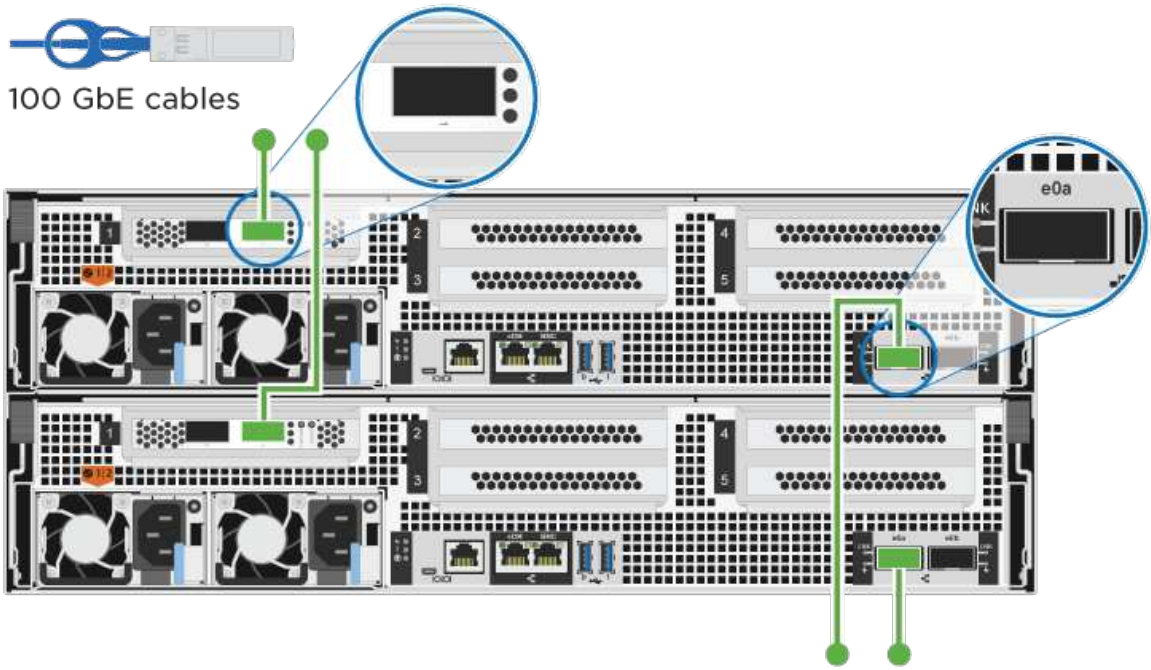
3	Cable the management ports to the management network switches 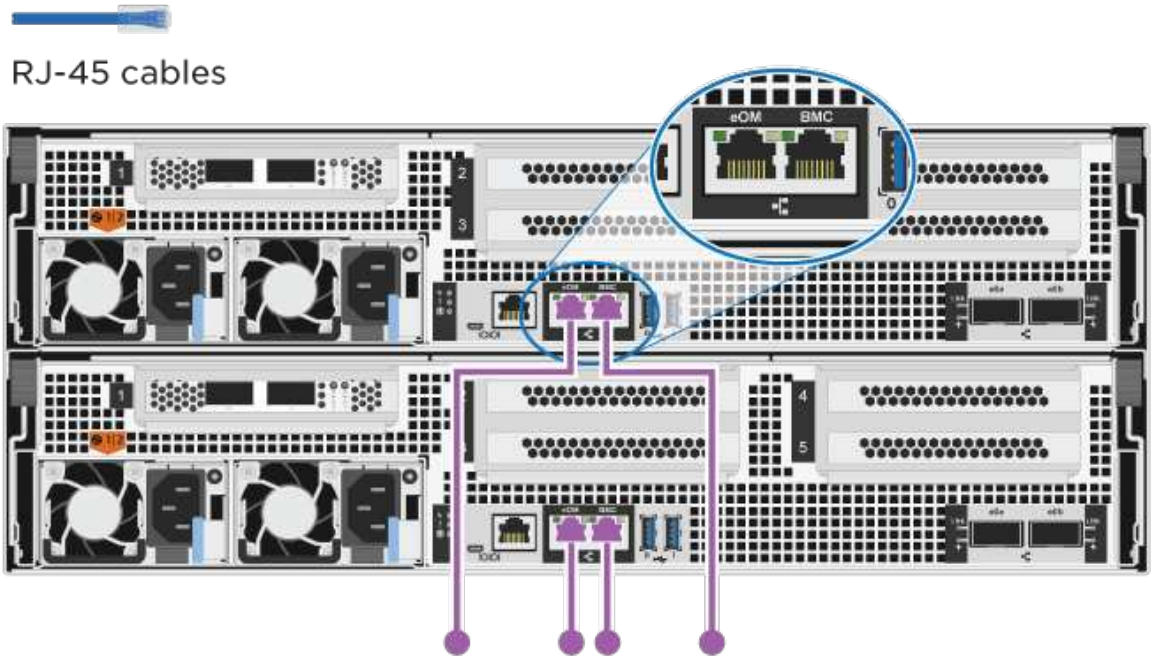

Step	Perform on each controller module
	DO NOT plug in the power cords at this point.

2. To perform optional cabling, see:
   - Option 1: Cable to a Fibre Channel host network
   - Option 2: Cable to a 10GbE host network
   - Option 3: Cable the controllers to a single drive shelf
   - Option 4: Cable the controllers to two drive shelves
3. To complete setting up your system, see Step 4: Complete system setup and configuration.

Optional cabling: Cable configuration-dependent options

You have configuration-dependent optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cabling to a host network and storage.

Option 1: Cable to a Fibre Channel host network

Fibre Channel ports on the controllers are connected to Fibre Channel host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

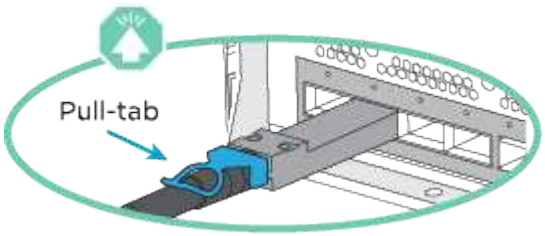

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Step	Perform on each controller module
1	Cable ports 2a through 2d to the FC host switches. 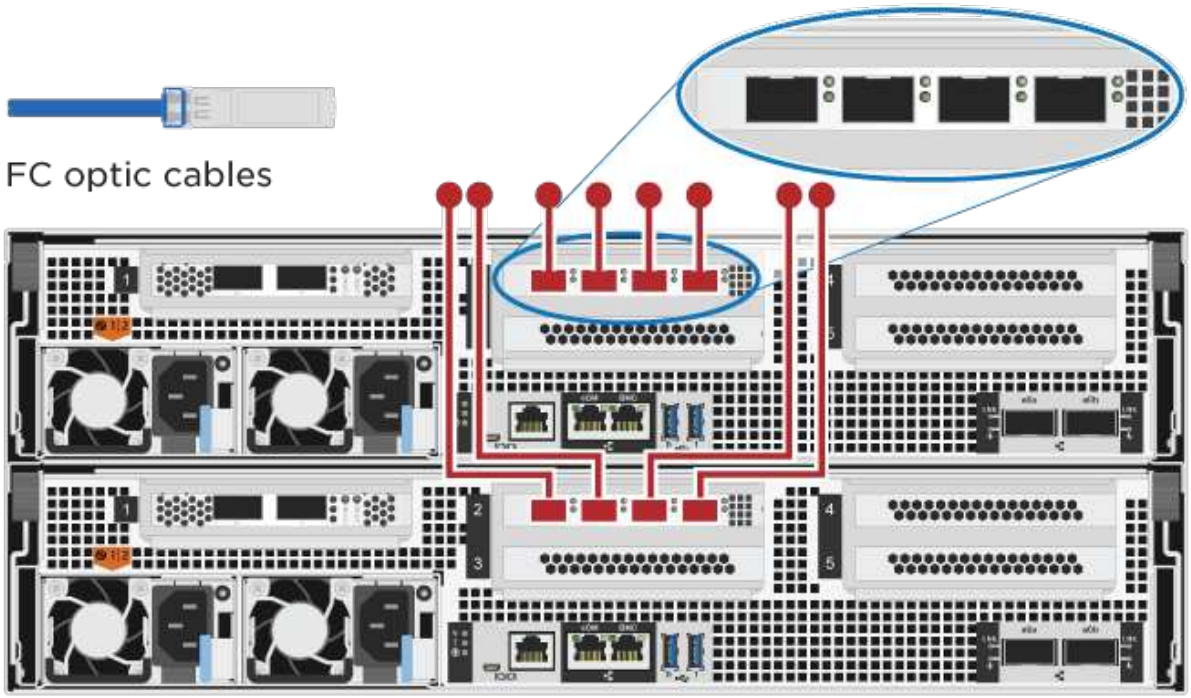
2	To perform other optional cabling, choose from: • Option 3: Cable the controllers to a single drive shelf • Option 4: Cable the controllers to two drive shelves
3	To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 2: Cable to a 10GbE host network

10GbE ports on the controllers are connected to 10GbE host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

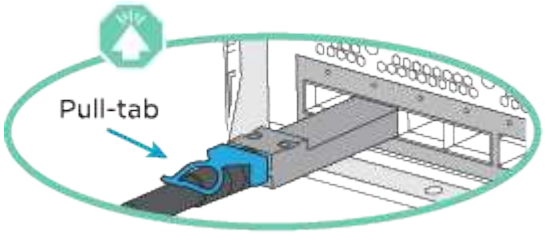

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Step	Perform on each controller module
1	Cable ports e4a through e4d to the 10GbE host network switches. 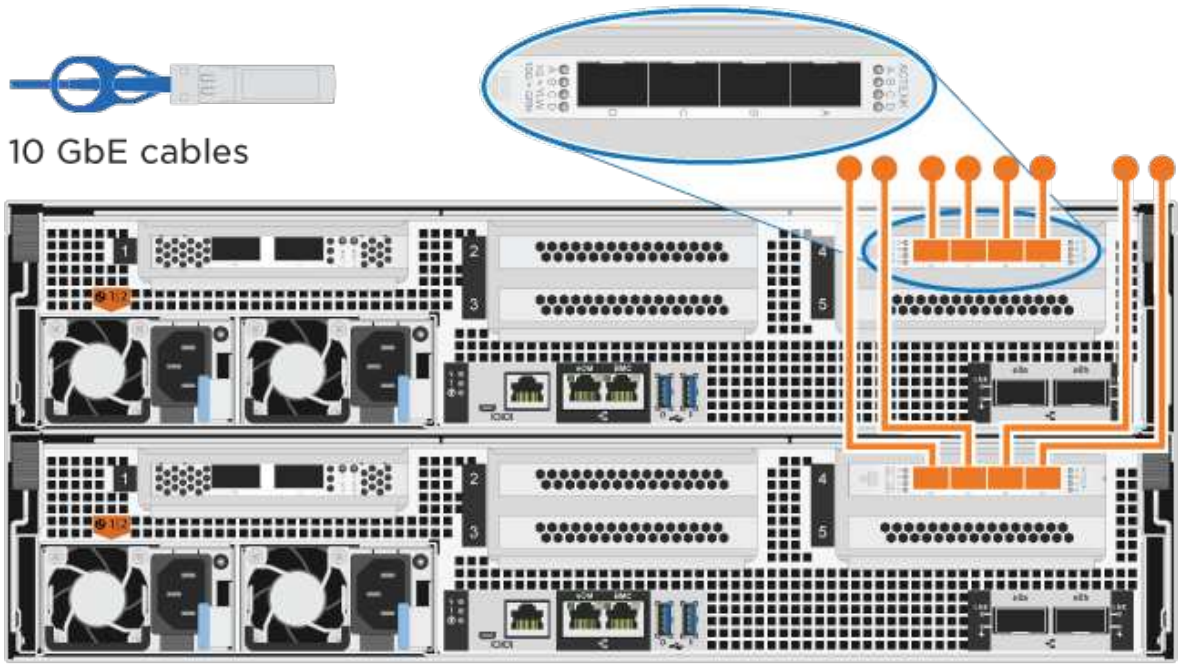
2	To perform other optional cabling, choose from: • Option 3: Cable the controllers to a single drive shelf • Option 4: Cable the controllers to two drive shelves
3	To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 3: Cable the controllers to a single drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

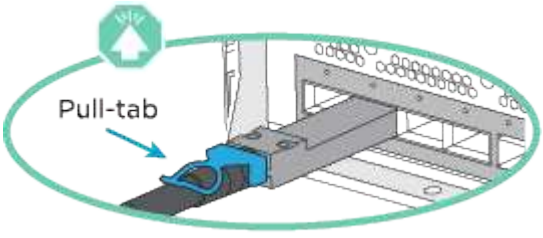

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to cable your controllers to a single shelf:

Animation - Cable the controllers to a single drive shelf

Step	Perform on each controller module
1	Cable controller A to the shelf: 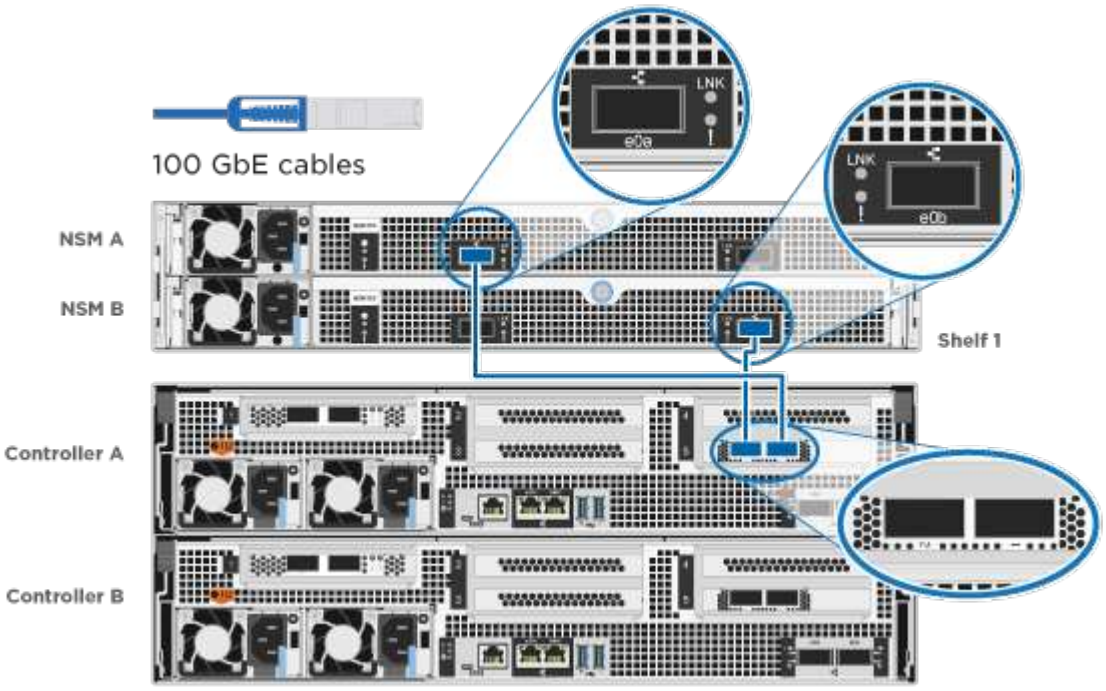
2	Cable controller B to the shelf: 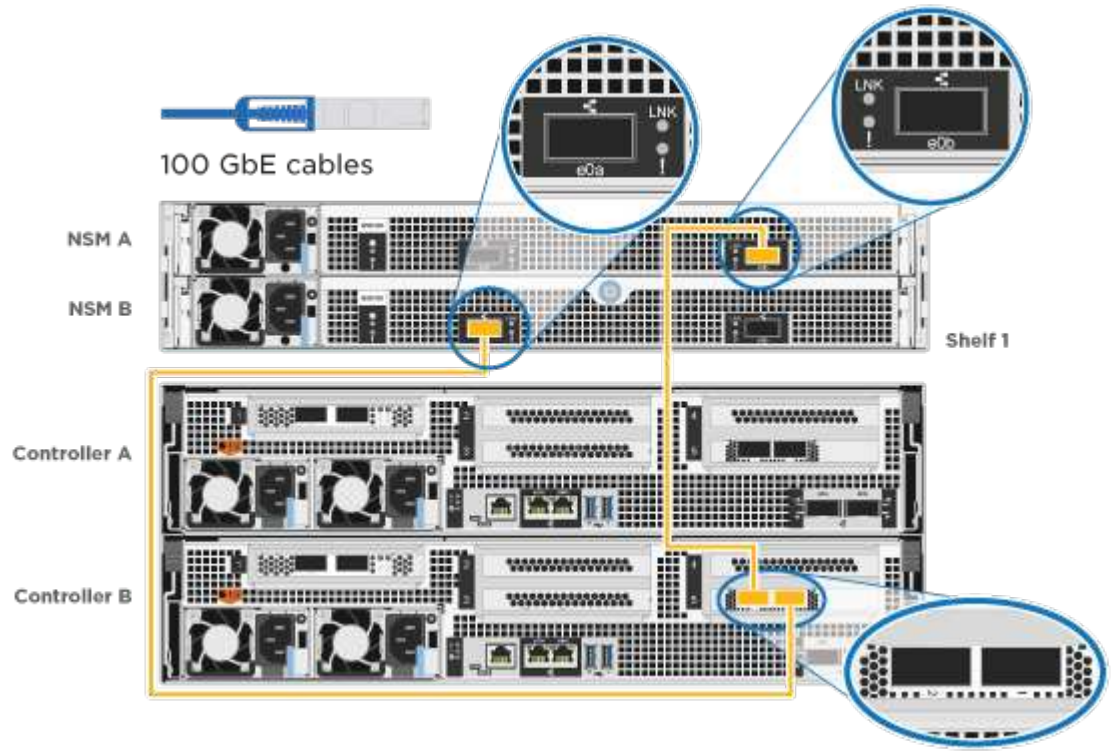

To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 4: Cable the controllers to two drive shelves

You must cable each controller to the NSM modules on both NS224 drive shelves.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

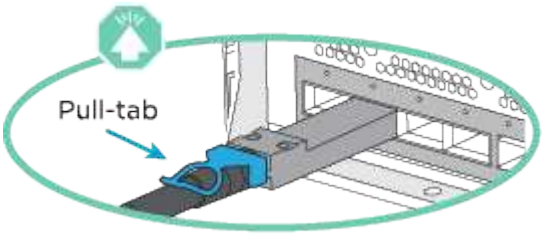

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to cable your controllers to two drive shelves:

Animation - Cable the controllers to two drive shelves

Step	Perform on each controller module
1	Cable controller A to the shelves: 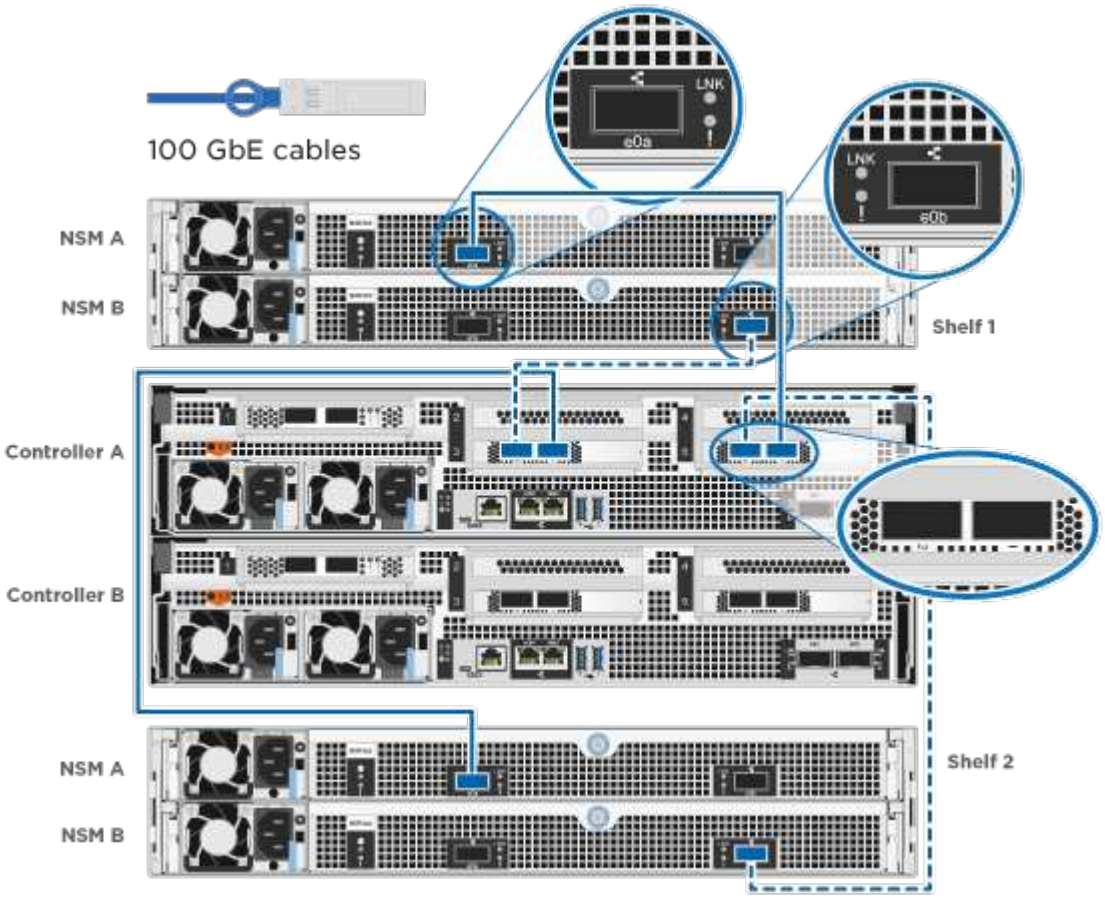

Step	Perform on each controller module
2	Cable controller B to the shelves: 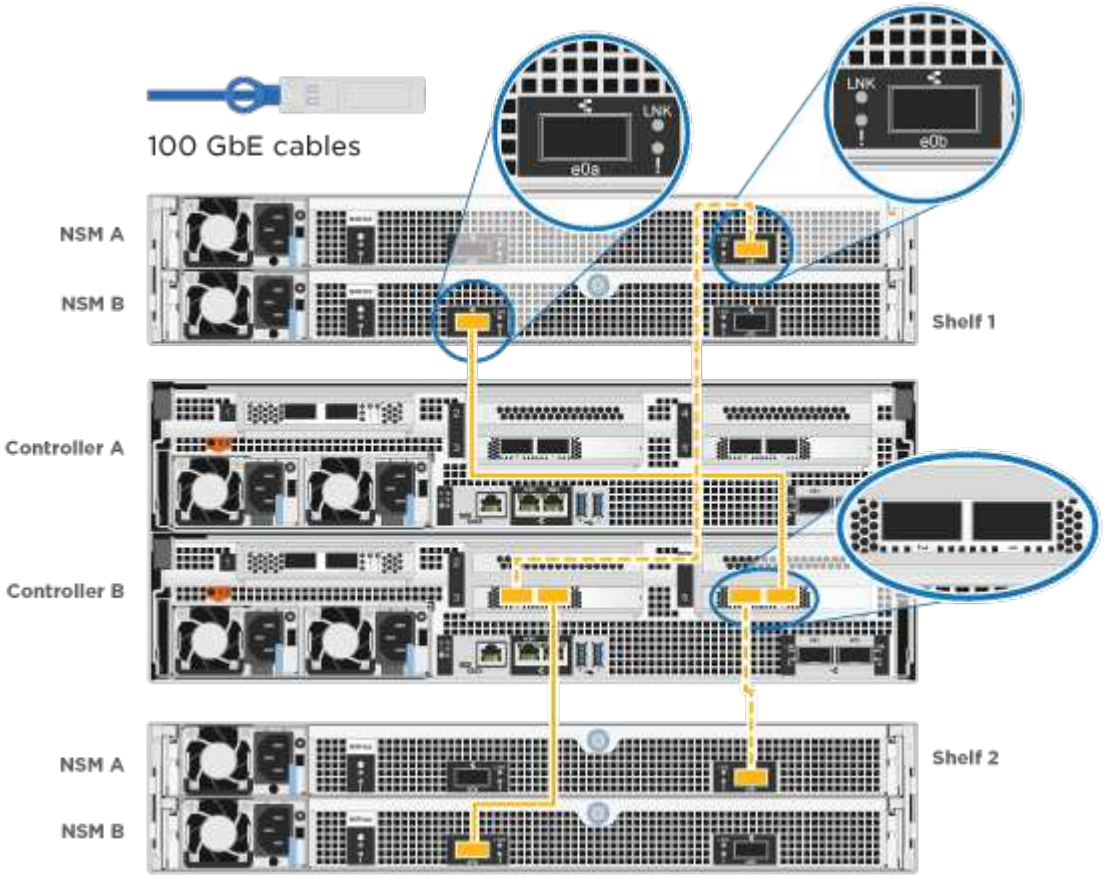

To complete setting up your system, see Step 4: Complete system setup and configuration.

Step 4: Complete system setup and configuration

Complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Complete system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

2. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

3. Use the animation to connect your laptop to the Management switch:

Animation - Connect your laptop to the Management switch

4. Select an ONTAP icon listed to discover:

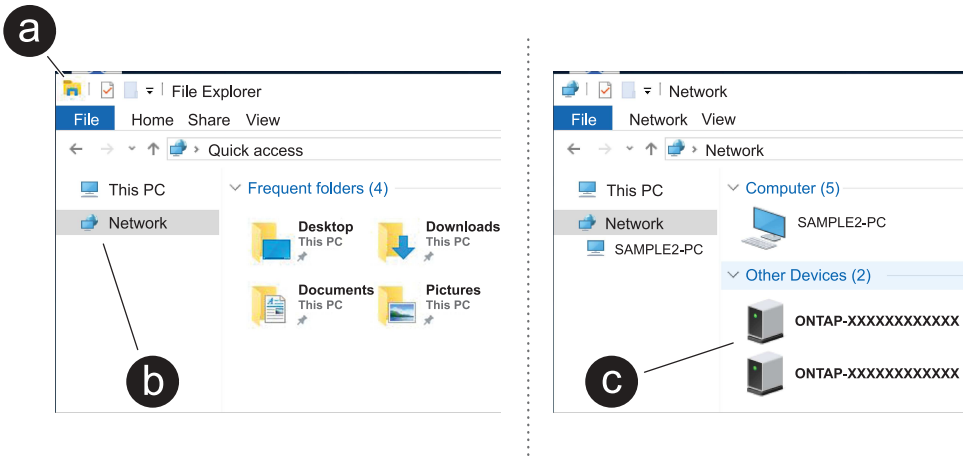

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

      System Manager opens.

5. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
6. Verify the health of your system by running Config Advisor.
7. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Complete system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your system.

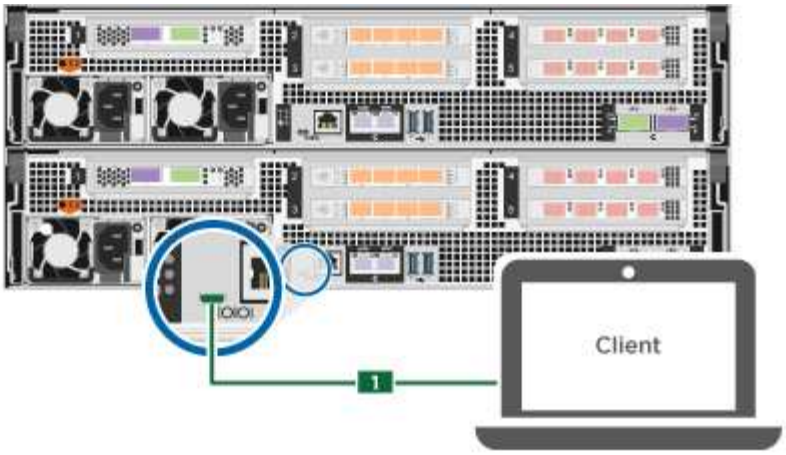

c. Connect the laptop or console to the switch on the management subnet.

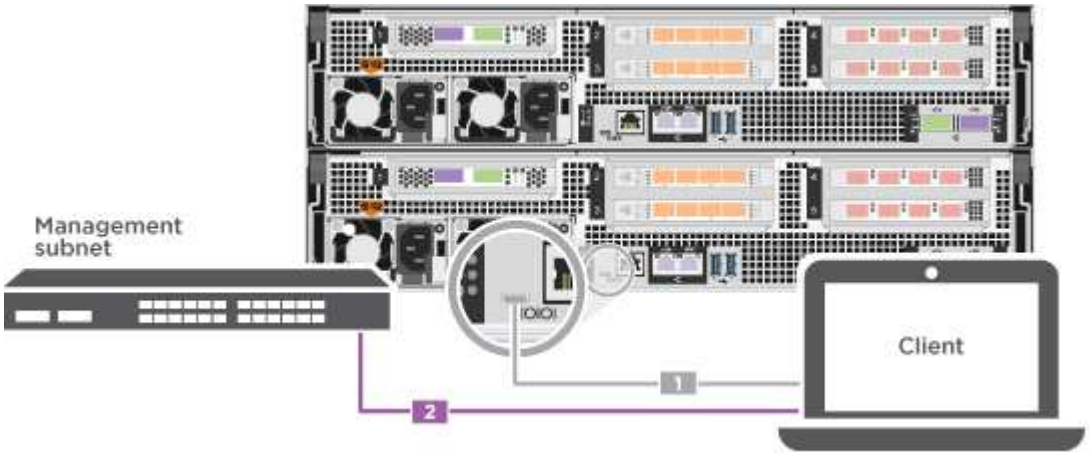

d. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

3. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

4. Using System Manager on your laptop or console, configure your cluster:

a. Point your browser to the node management IP address.

> The format for the address is https://x.x.x.x.

b. Configure the system using the data you collected in the ONTAP Configuration Guide.

5. Verify the health of your system by running Config Advisor.
6. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain ASA A800 hardware

Maintain the hardware of your ASA A800 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the ASA A800 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the ASA A800 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Drive	A drive is a device that provides the physical storage media for data.

Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.
NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.
PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - ASA A800

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your ASA A800 system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5 **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - ASA A800

Before replacing the boot media in your ASA A800, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - ASA A800

Shut down the impaired controller in your ASA A800 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - ASA A800

The boot media in your ASA A800 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module,

removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

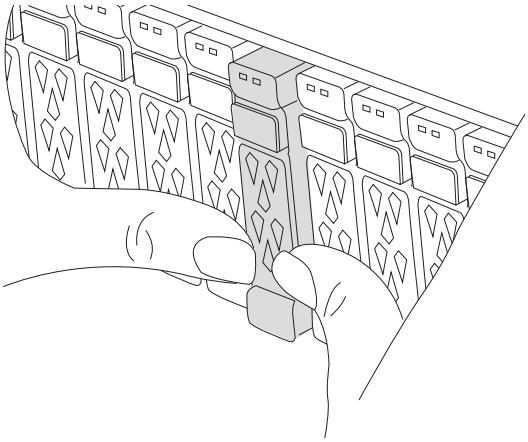

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

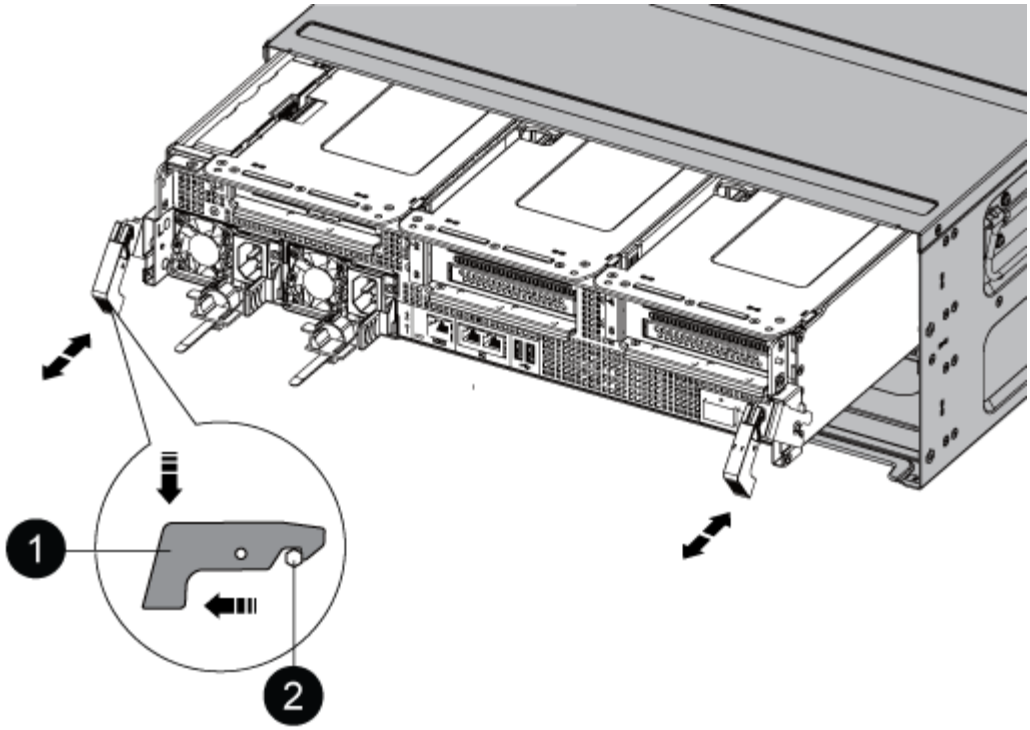

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

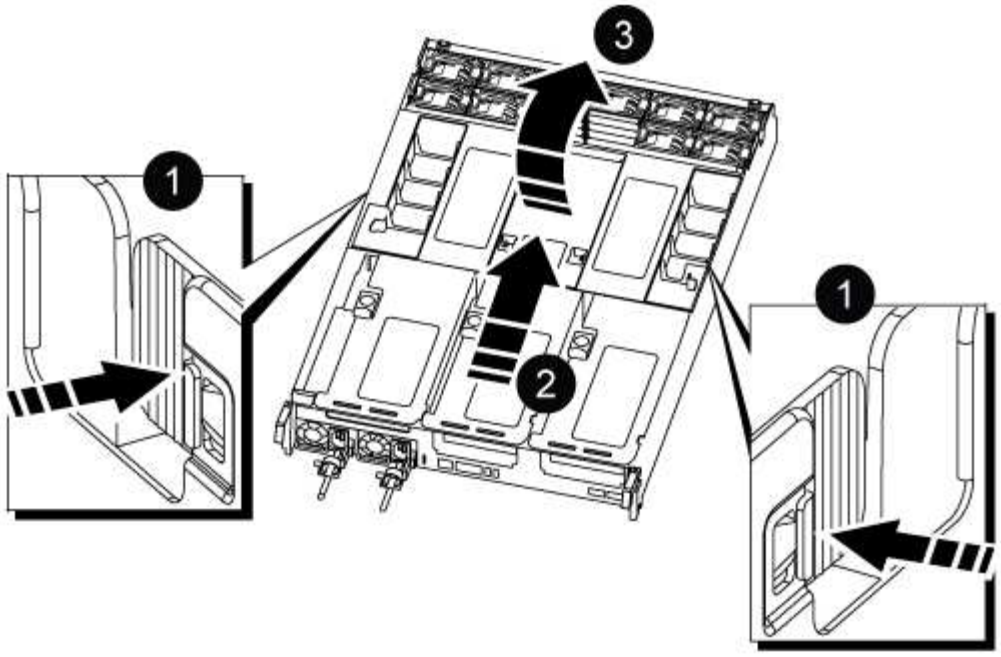

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

10. Locate the boot media in the controller module and replace it:

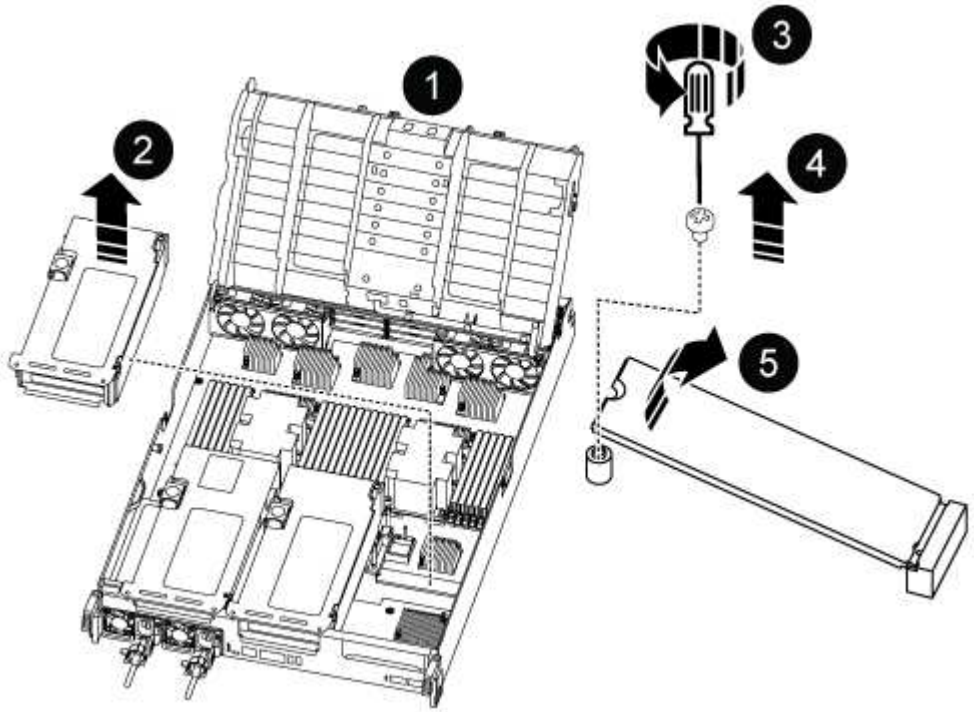

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.

b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.

11. Install the replacement boot media into the controller module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down toward the motherboard.

c. Secure the boot media to the motherboard using the boot media screw.

Do not over-tighten the screw or you might damage the boot media.

12. Reinstall the riser into the controller module.

13. Close the air duct:

a. Rotate the air duct downward.

b. Slide the air duct toward the risers until it clicks into place.

14. Install the controller module:

a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module half-way into the way into the system.

b. Recable the controller module, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

The controller module begins to boot and stops at the LOADER prompt.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - ASA A800

After installing the new boot media device in your ASA A800 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and

determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

   Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

   Show example of configuration error finding prompts

   ```
   Error when fetching key manager config from partner ${partner_ip}:
   ${status}

   Has key manager been configured on this system

   Is the key manager onboard
   ```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - ASA A800

If a component in your ASAF A800 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - ASA A800

Get started with replacing the boot media in your ASA A800 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - ASA A800

Before replacing the boot media in your ASA A800 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - ASA A800

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - ASA A800

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: System is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media and prepare for manual boot recovery - ASA A800

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

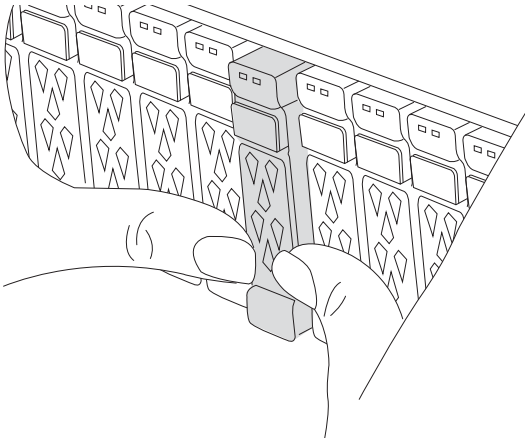

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management

device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

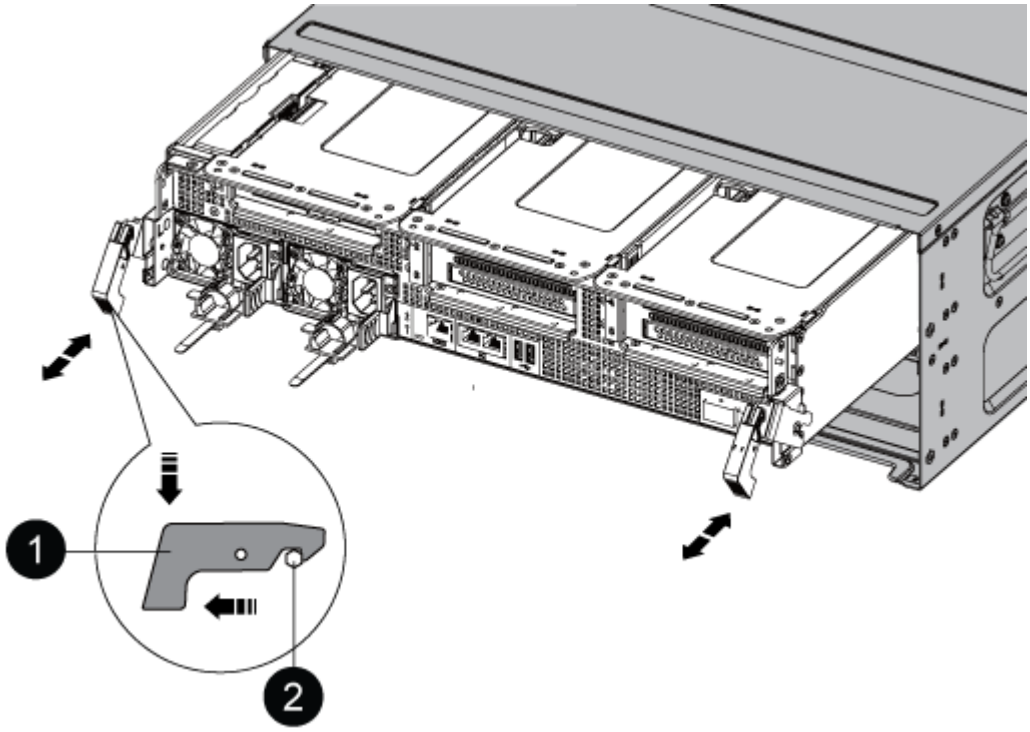

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

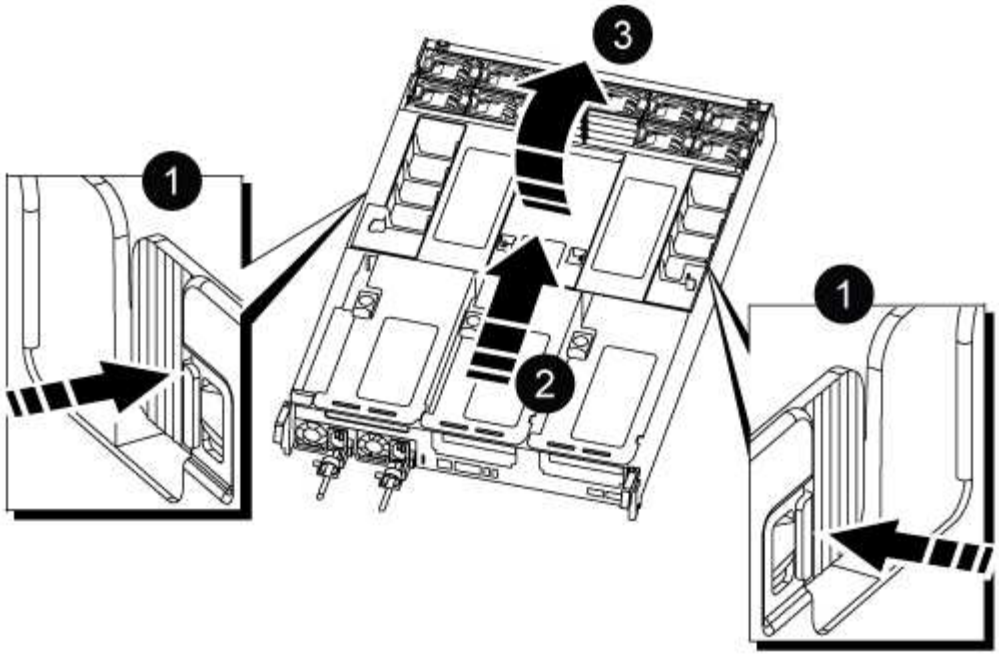

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 2: Replace the boot media

You locate the failed boot media in the controller module by removing Riser 3 on the controller module before you can replace the boot media.

You need a Phillips head screwdriver to remove the screw that holds the boot media in place.

1. Locate the boot media:

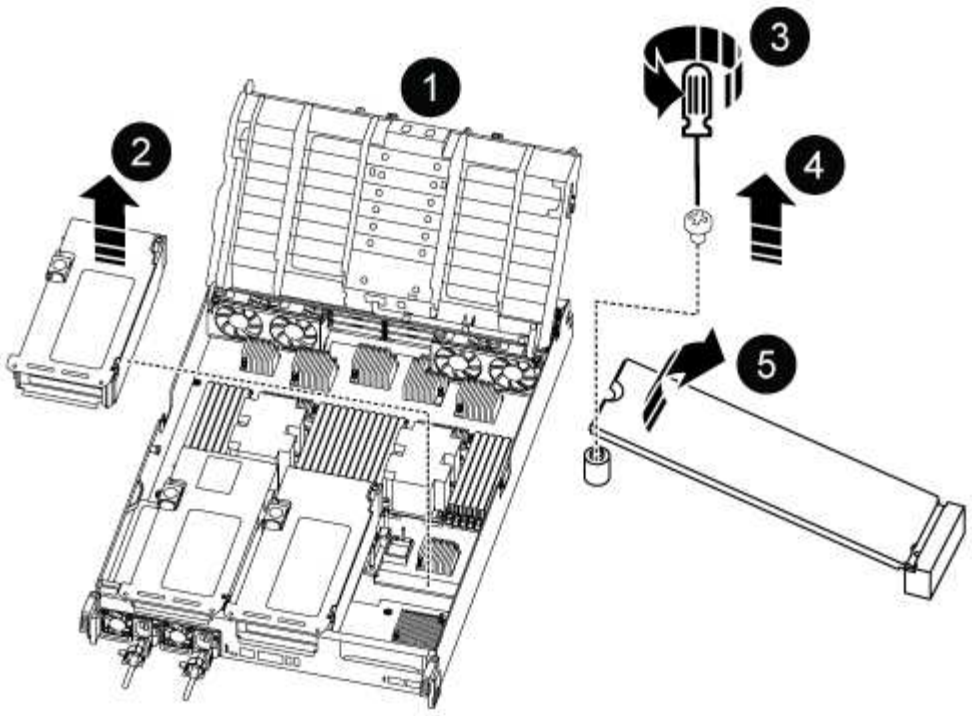

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

2. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.
   b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.
3. Install the replacement boot media into the controller module:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the motherboard.
   c. Secure the boot media to the motherboard using the boot media screw.

      Do not over-tighten the screw or you might damage the boot media.
4. Reinstall the riser into the controller module.

5. Close the air duct:
    a. Rotate the air duct downward.
    b. Slide the air duct toward the risers until it clicks into place.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed is without a boot image so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
    - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
    - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
    a. Download the service image to your work space on your laptop.
    b. Unzip the service image.

        If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

        There are two folders in the unzipped service image file:

        - boot
        - efi

    c. Copy the efi folder to the top directory on the USB flash drive.

        If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

    The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

    a. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct:
    a. Swing the air duct all the way down to the controller module.
    b. Slide the air duct toward the risers until the locking tabs click into place.

c. Inspect the air duct to make sure that it is properly seated and locked into place.

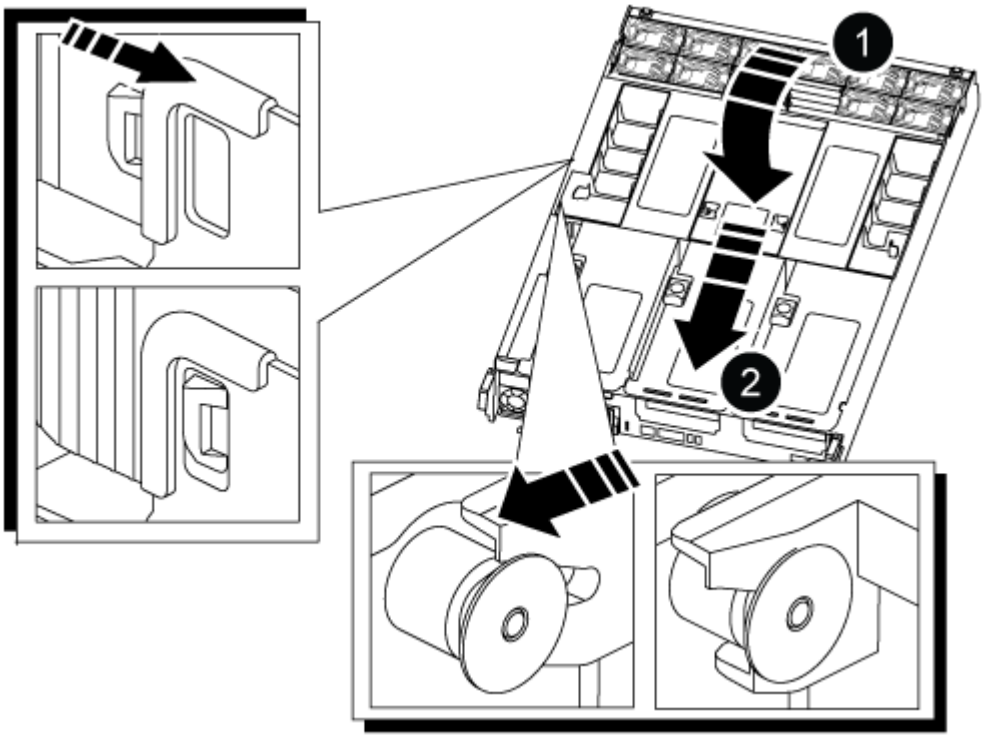

1	Air duct
2	Risers

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.

5. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

6. Gently push the controller module all the way into the system until the controller module locking hooks begin to rise, firmly push on the locking hooks to finish seating the controller module, and then swing the locking hooks into the locked position over the pins on the controller module.
7. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the

controller to boot to LOADER.

Manual boot media recovery from a USB drive - ASA A800

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore OKM, NSE, and NVE - ASA A800

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - ASA A800

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Replace the chassis - ASA A800

To replace the chassis, you must move the bezel, controller modules, and NVMe drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in its place.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - ASA A800

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Replace chassis - ASA A800

Move the power supplies, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

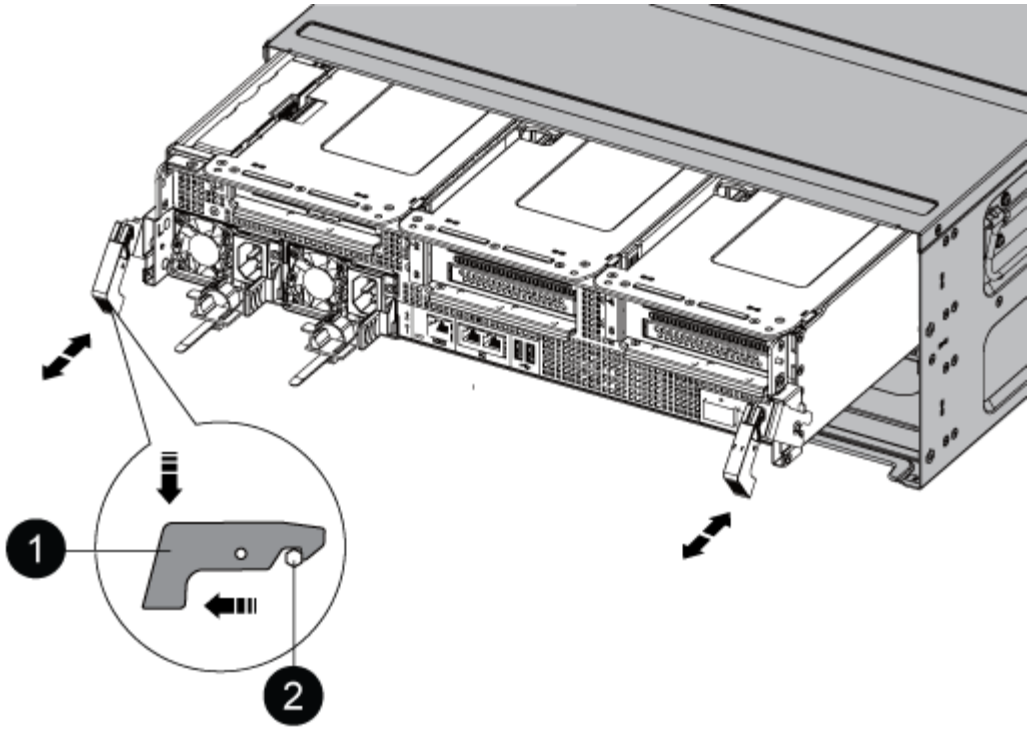

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.

3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.

5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process by pressing `Ctrl-C`.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete chassis replacement - ASA A800

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller replacement - ASA A800

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Do not downgrade the BIOS version of the *replacement* controller to match the partner controller or the old controller module.

Shut down the impaired controller - ASA A800

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   `cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller - ASA A800

To replace the controller, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

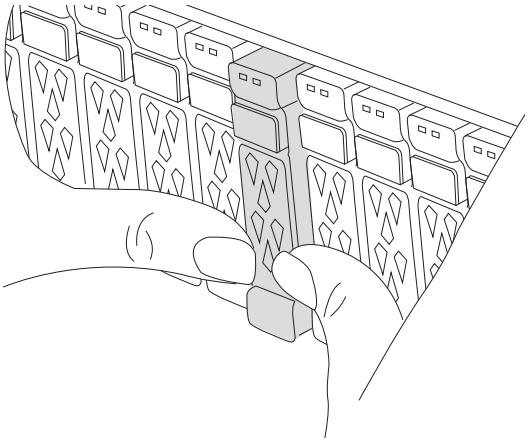

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

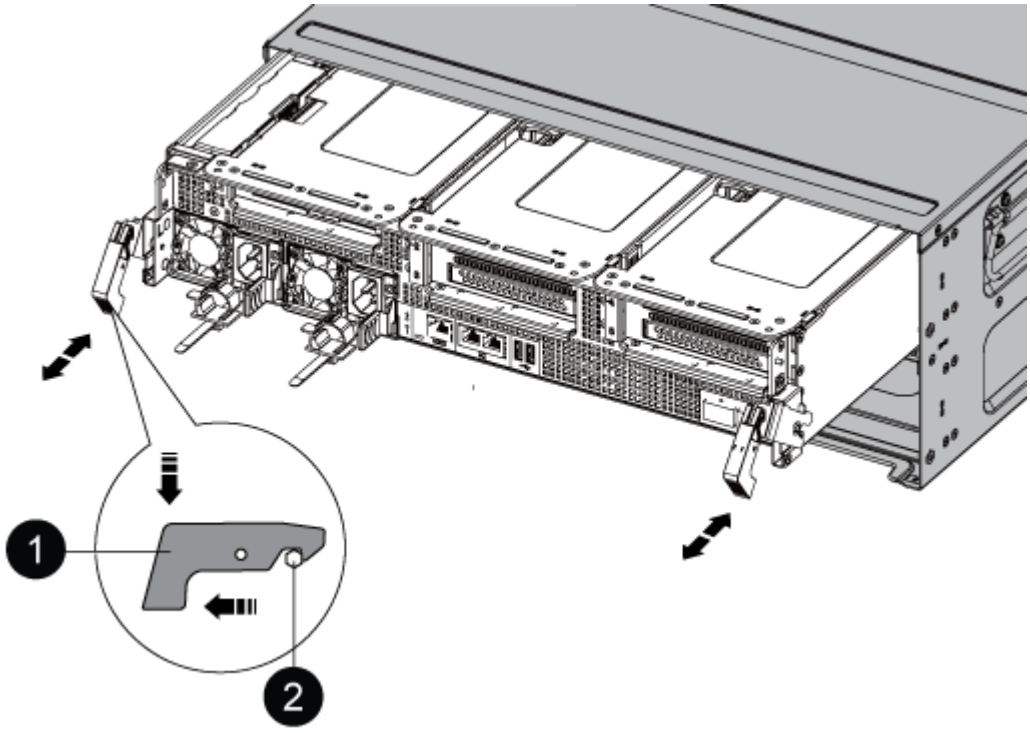

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
   .

9. Place the controller module on a stable, flat surface.
10. Open the controller module air duct:
    a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
    b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

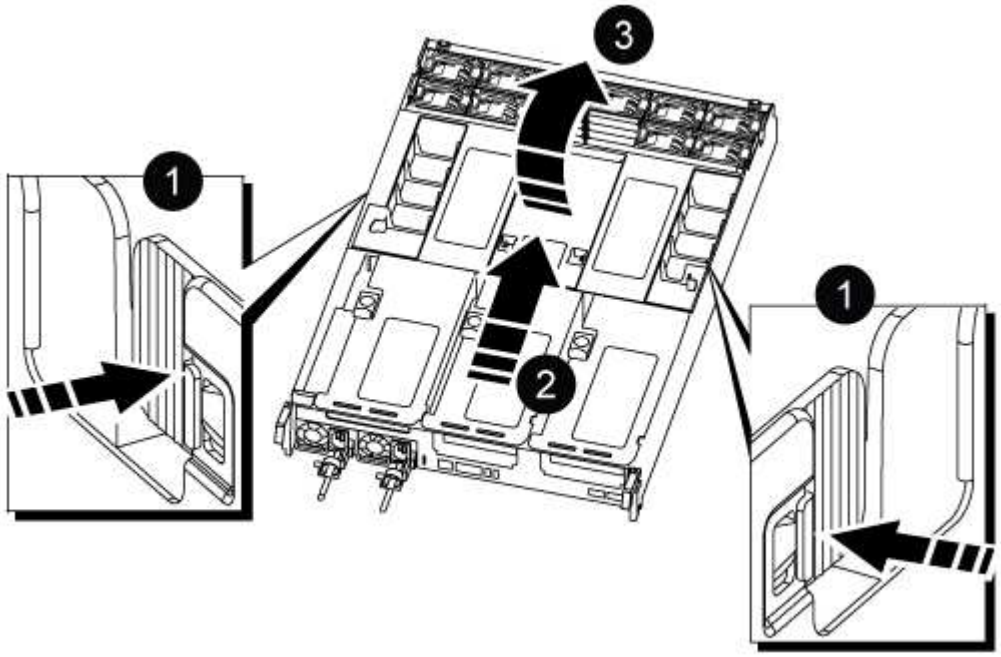

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 2: Move the power supplies

You must move the power supplies from the impaired controller module to the replacement controller module when you replace a controller module.

1. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

> The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

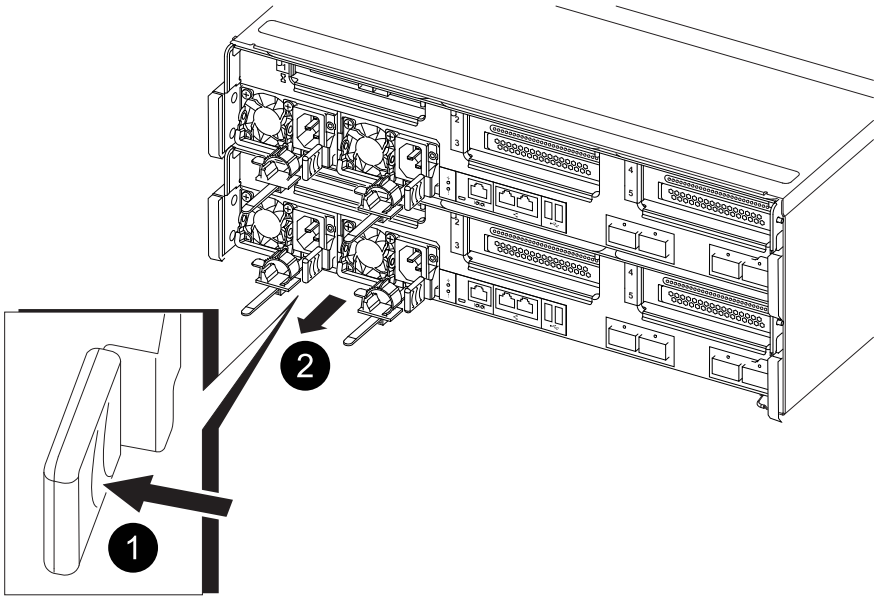

1	Blue power supply locking tab
2	Power supply

2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

> To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

You must move the fans from the impaired controller module to the replacement module when replacing a failed controller module.

1. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

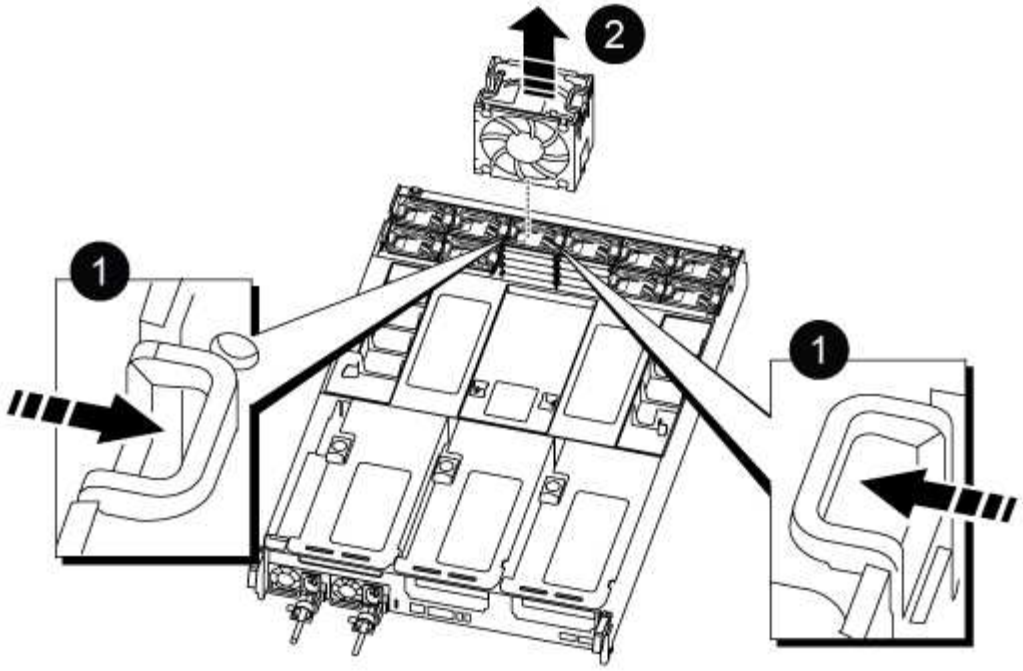

1	Fan locking tabs
2	Fan module

2. Move the fan module to the replacement controller module, and then install the fan module by aligning its edges with the opening in the controller module, and then sliding the fan module into the controller module until the locking latches click into place.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the NVDIMM battery

When replacing the controller module, you must move the NVRAM battery from the impaired controller module to the replacement controller module

1. Open the air duct cover and locate the NVDIMM battery in the riser.

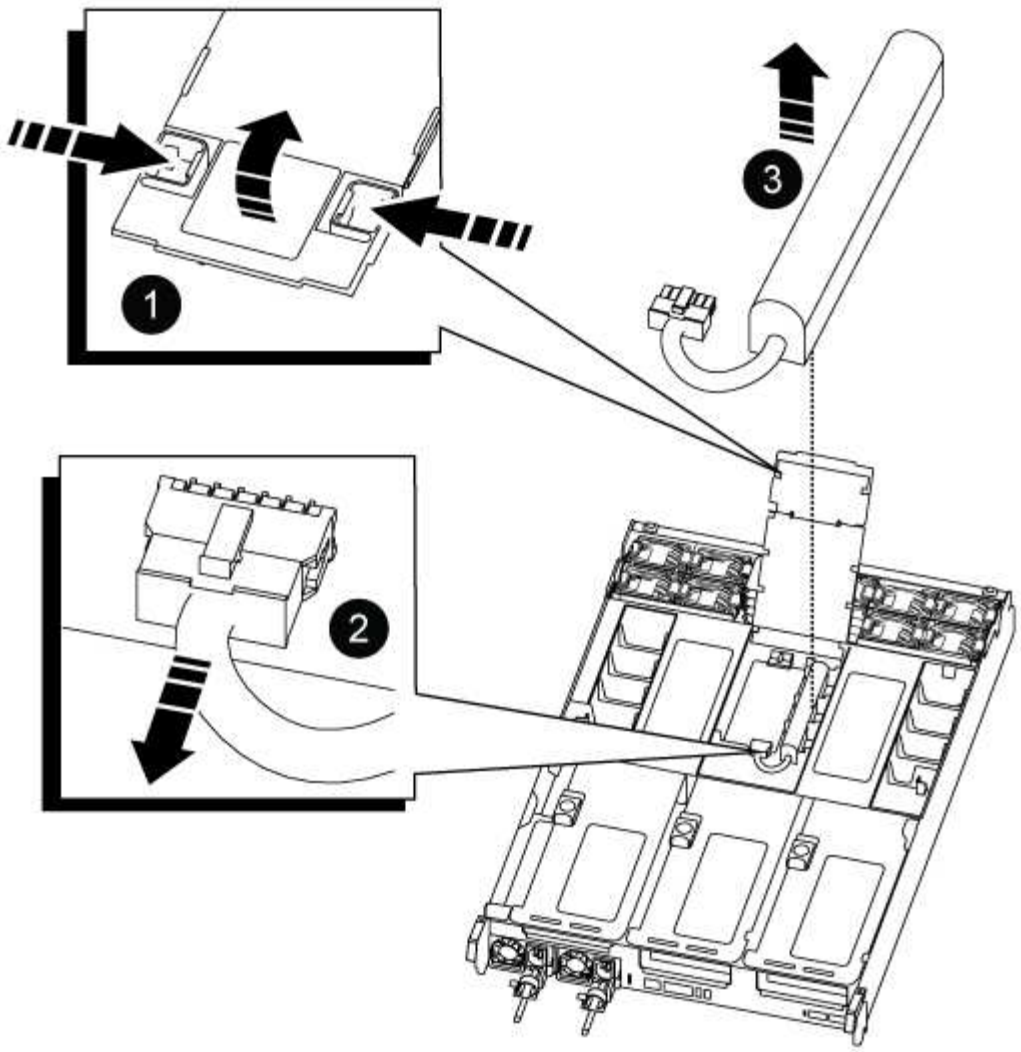

1	Air duct riser
2	NVDIMM battery plug
3	NVDIMM battery pack

Attention: The NVDIMM battery control board LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
3. Grasp the battery and lift the battery out of the air duct and controller module.
4. Move the battery pack to the replacement controller module and then install it in the NVDIMM air duct:
   a. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   b. Plug the battery plug into the riser socket and make sure that the plug locks into place.

Step 5: Remove the PCIe risers

As part of the controller replacement process, you must remove the PCIe modules from the impaired controller module. You must install them into the same location in the replacement controller module once the NVDIMMS and DIMMs have moved to the replacement controller module.

1. Remove the PCIe riser from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

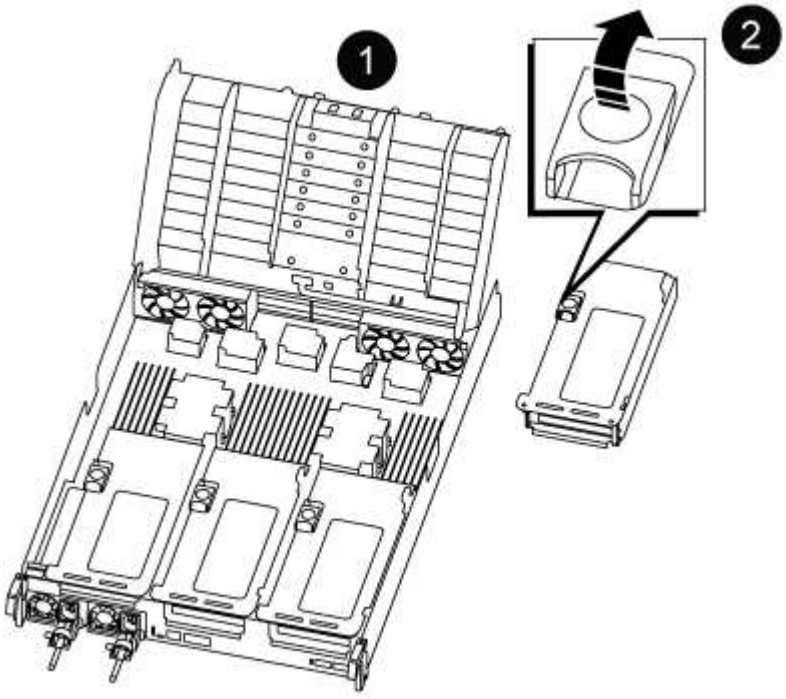

1	Air duct
2	Riser 1 (left riser), Riser 2 (middle riser), and 3 (right riser) locking latches

2. Repeat the preceding step for the remaining risers in the impaired controller module.
3. Repeat the above steps with the empty risers in the replacement controller and put them away.

Step 6: Move system DIMMs

To move the DIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

1. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.

2. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

3. Locate the slot where you are installing the DIMM.
4. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

5. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
6. Repeat these steps for the remaining DIMMs.

Step 7: Move the NVDIMMs

To move the NVDIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

1. Locate the NVDIMMs on your controller module.

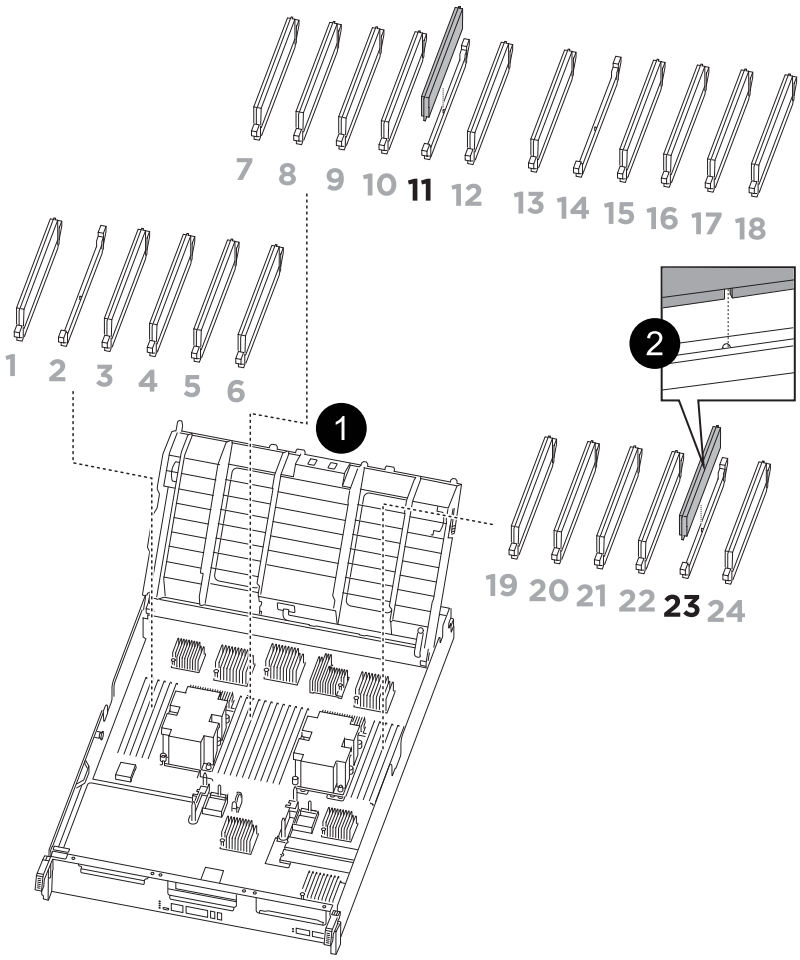

- NVDIMM: SLOTS 11 & 23

1	Air duct
2	NVDIMMs

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

4. Locate the slot where you are installing the NVDIMM.
5. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
7. Repeat the preceding steps to move the other NVDIMM.

Step 8: Move the boot media

You must move the boot media device from the impaired controller and install it in the replacement controller.

The boot media is located under Riser 3.

1. Locate the boot media:

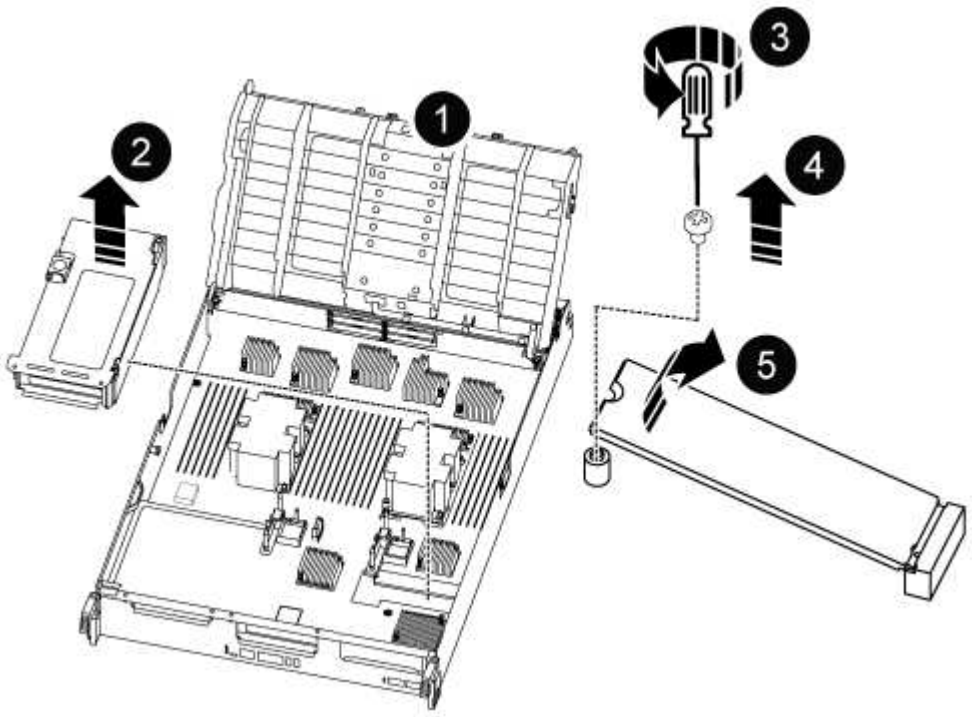

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

2. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.

b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.

3. Move the boot media to the new controller module and install it:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the motherboard.
   c. Secure the boot media to the motherboard using the boot media screw.

      Do not over-tighten the screw or you might damage the boot media.

Step 9: Install the PCIe risers

You install the PCIe risers in the replacement controller module after moving the DIMMs, NVDIMMs, and boot media.

1. Install the riser into the replacement controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

   d. Reinsert any SFP or QSFP modules that were removed from the PCIe cards.
2. Repeat the preceding step for the remaining PCIe risers.

Step 10: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

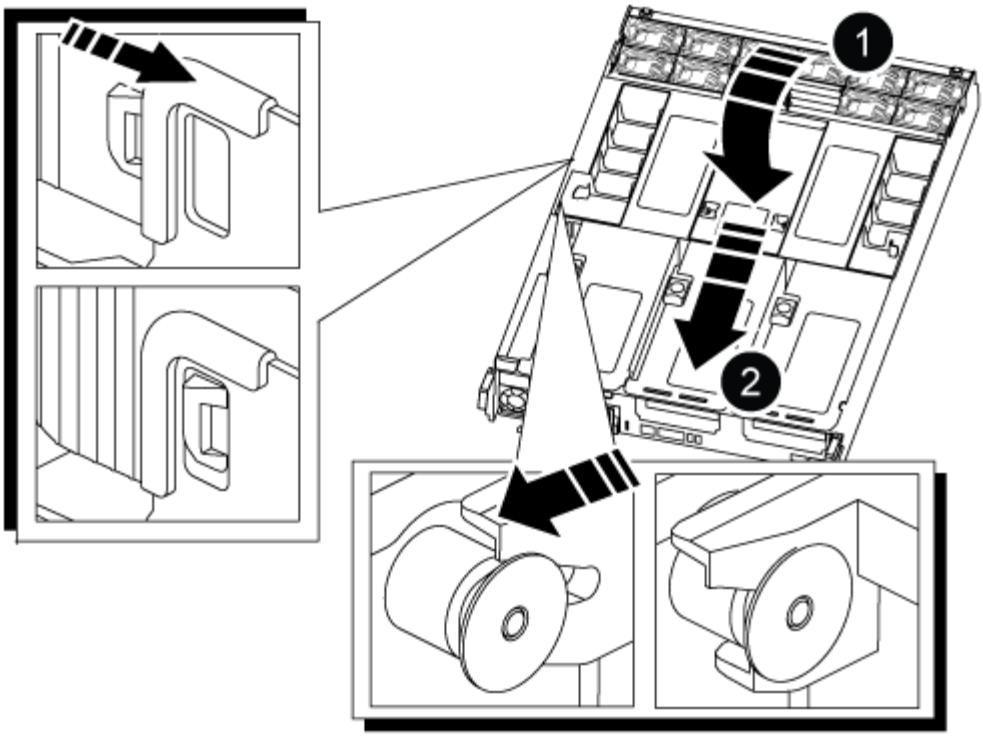

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:

   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

5. Plug the system cables and transceiver modules into the controller module and reinstall the cable management device.

6. Plug the power cables into the power supplies and reinstall the power cable retainers.

   The controller module begins to boot as soon as it is connected to power. Be prepared to interrupt the boot process.

   If your system has DC power supplies, make sure the thumbscrews on the power supply cable are tight.

Restore and verify the system configuration - ASA A800

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the chassis

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA`

state: `ha-config show`

The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - ASA A800

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and

then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - ASA A800

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace DIMM - ASA A800

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

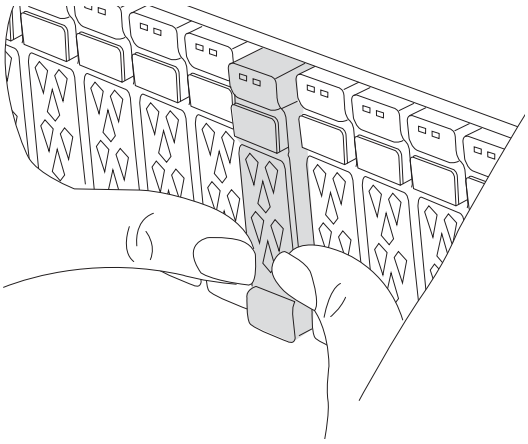

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

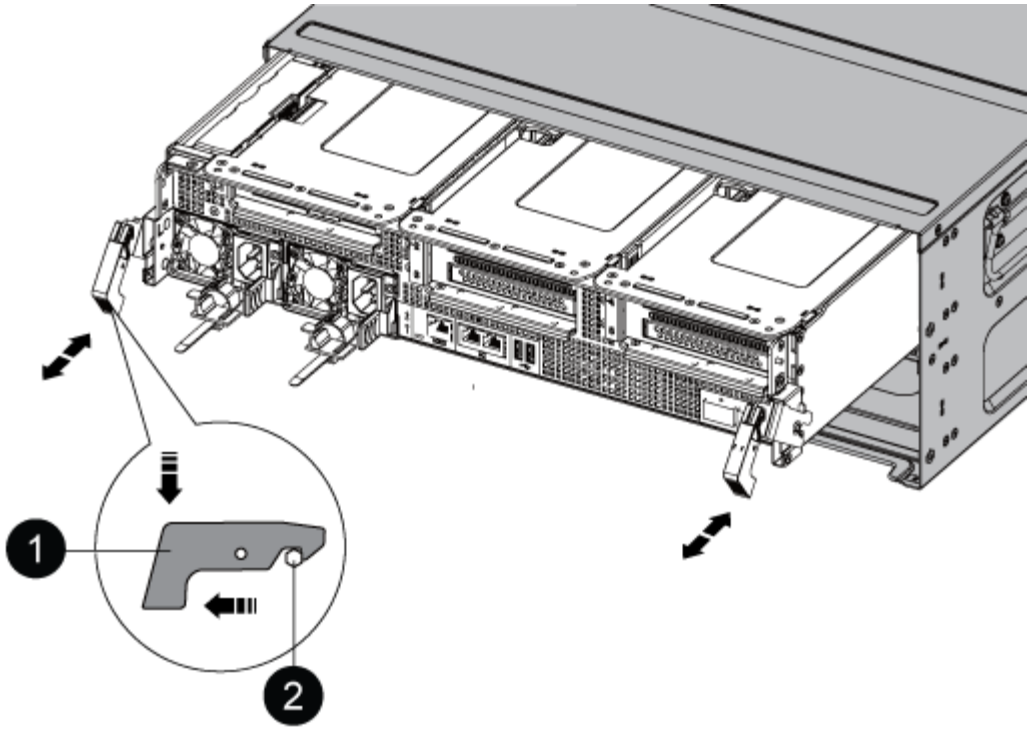

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

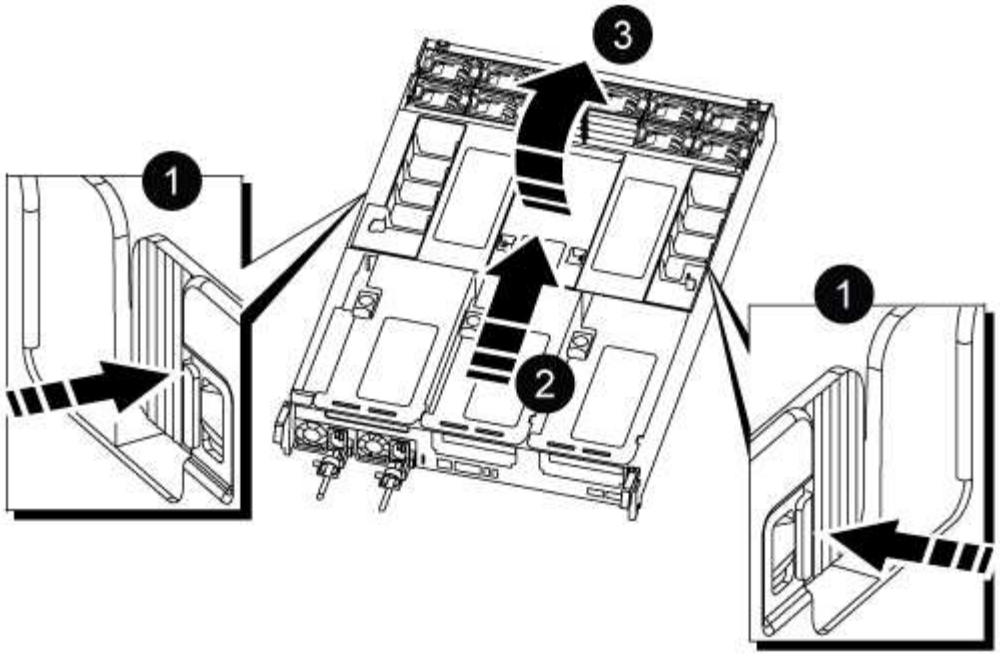

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace a DIMM

To replace a DIMM, you must locate it in the controller module using the DIMM map label on top of the air duct and then replace it following the specific sequence of steps.

1. When removing a DIMM, unlock the locking latch on the applicable riser, and then remove the riser.

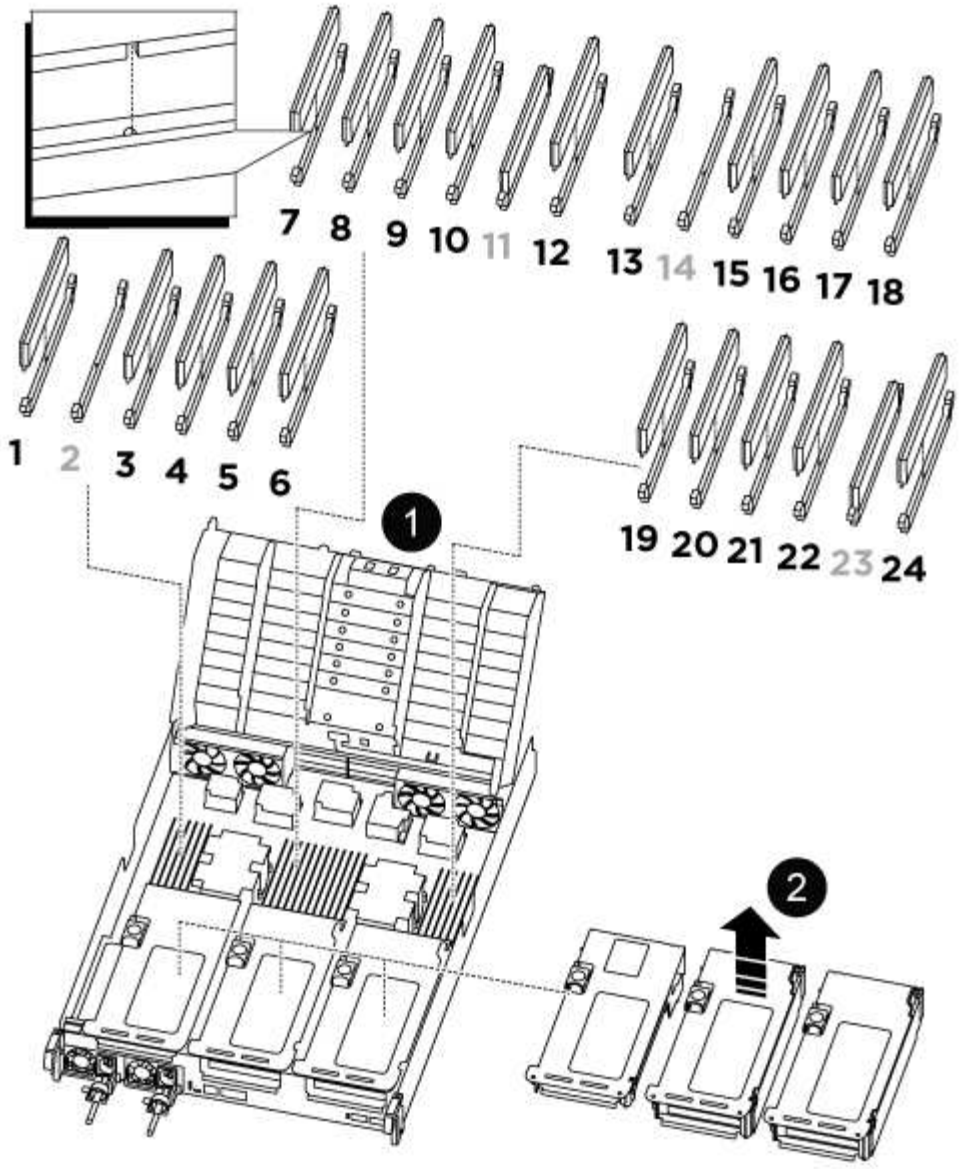

1	Air duct cover
2	Riser 1 and DIMM bank 1, and 3-6
Riser 2 and DIMM bank 7-10, 12-13, and 15-18	Riser 3 and DIMM 19 -22 and 24

Note: Slot 2 and 14 are left empty. Do not attempt to install DIMMs into these slots.

2. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

4. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

5. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Reinstall any risers that you removed from the controller module.
8. Close the air duct.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

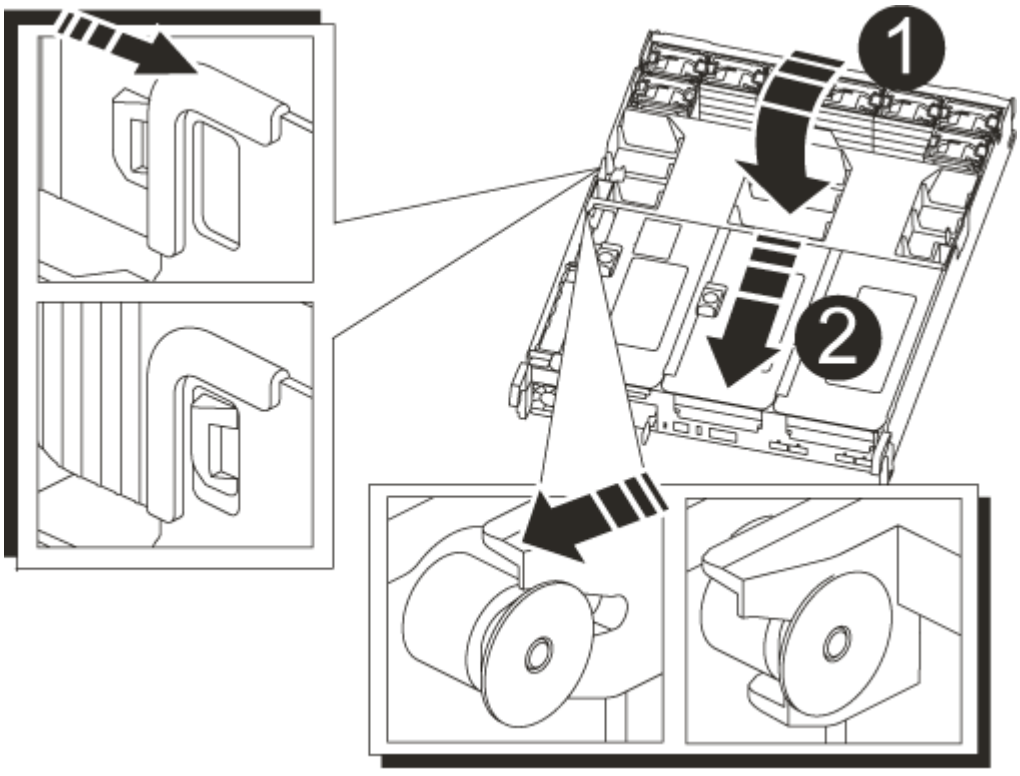

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - ASA A800

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before

replacing a drive.

- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  > Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   > The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan - ASA A800

To replace a fan, remove the failed fan module and replace it with a new fan module.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a fan module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

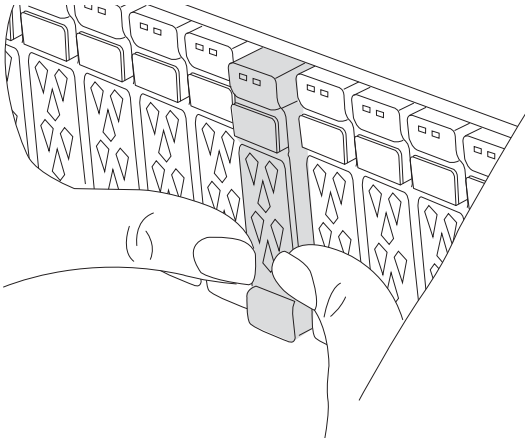

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

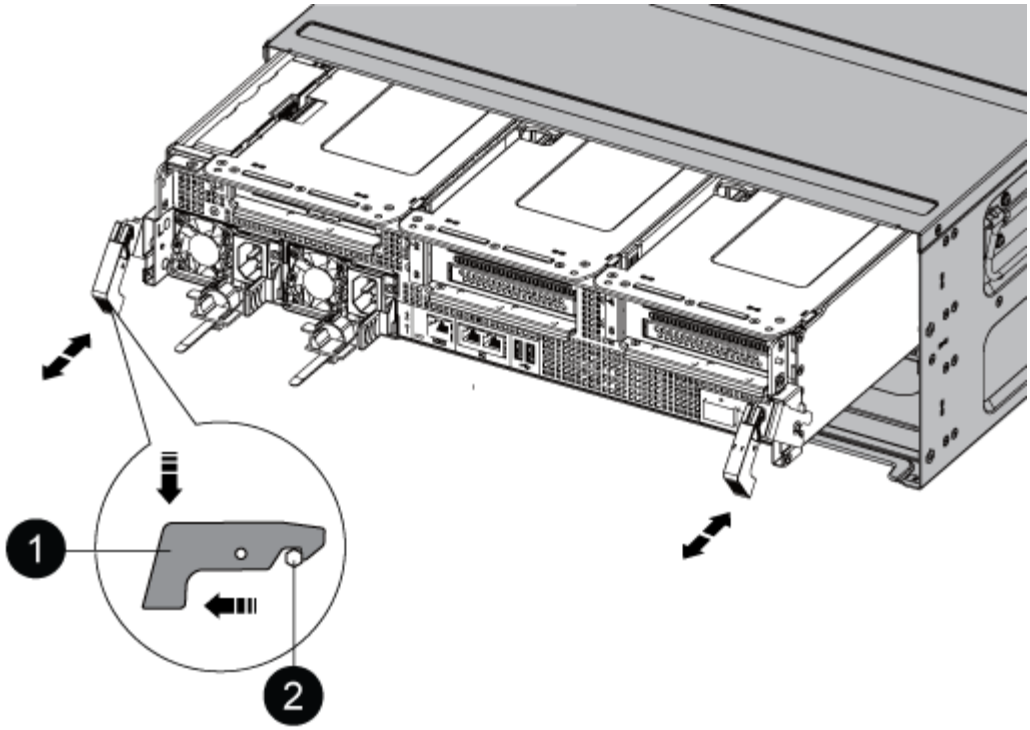

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Set the controller module aside in a safe place.

Step 3: Replace a fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
2. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

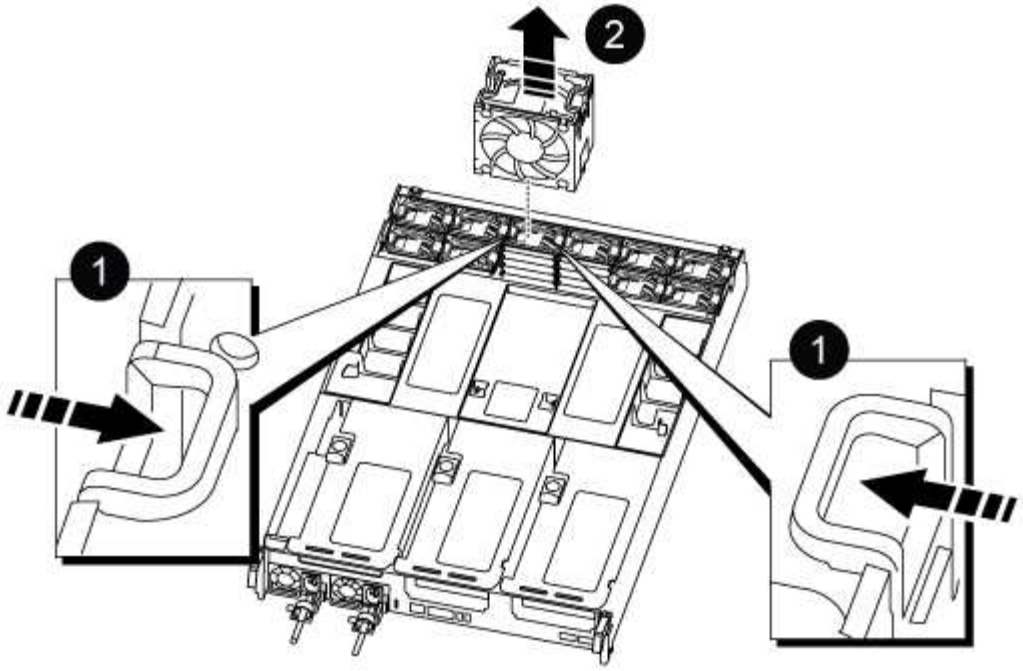

1	Fan locking tabs
2	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module until the locking latches click into place.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.
3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -controller local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an NVDIMM - ASA A800

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

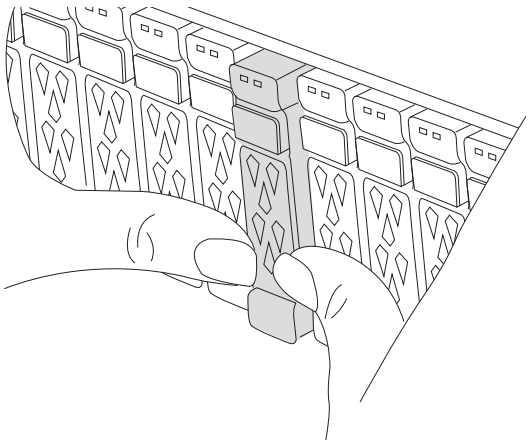

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.

5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

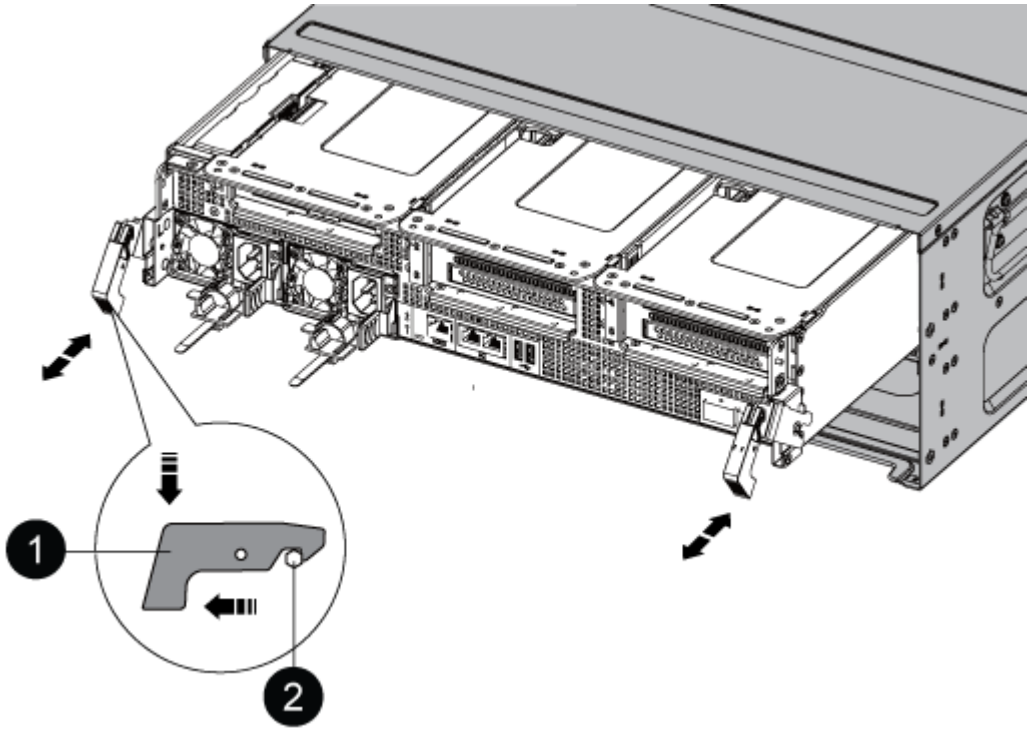

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:

   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

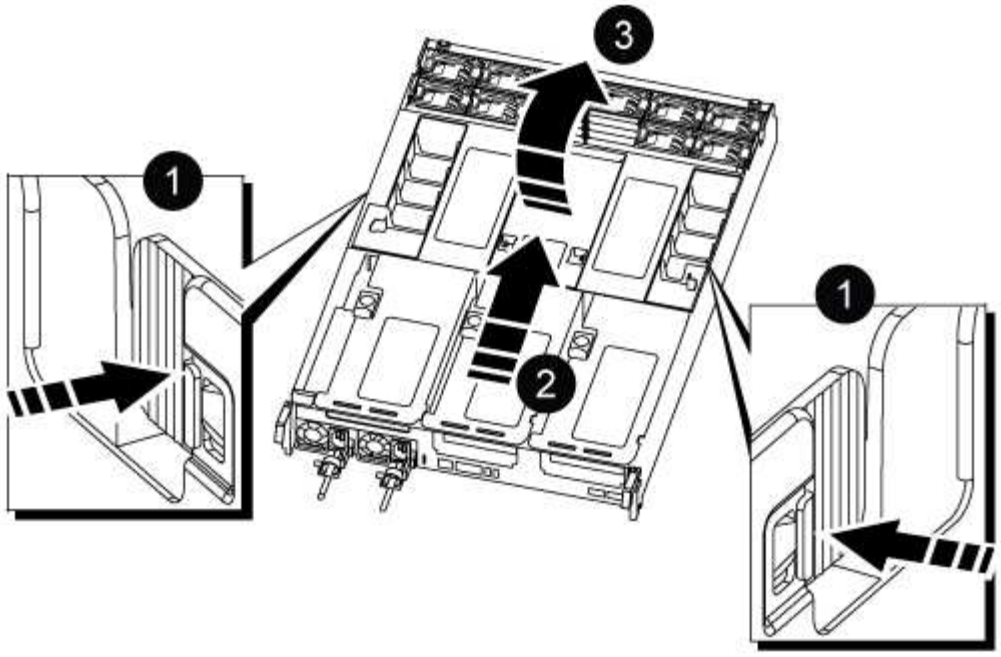

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the NVDIMM map label on top of the air duct, and then replace it following the specific sequence of steps.

1. If you are removing or moving an NVDIMM, unlock the locking latch on the riser, and then remove the applicable riser.

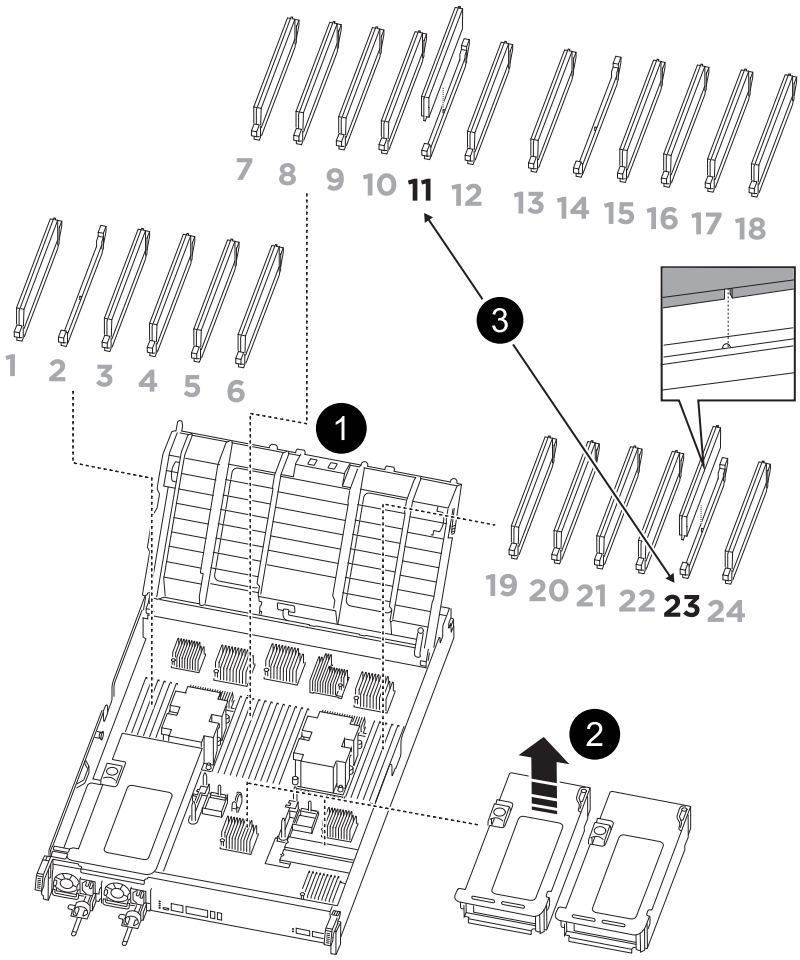

1	Air duct cover
2	Riser 2
3	NVDIMM in slots 11 and 23

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

4. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

5. Locate the slot where you are installing the NVDIMM.

6. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
8. Reinstall any risers that you removed from the controller module.
9. Close the air duct.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

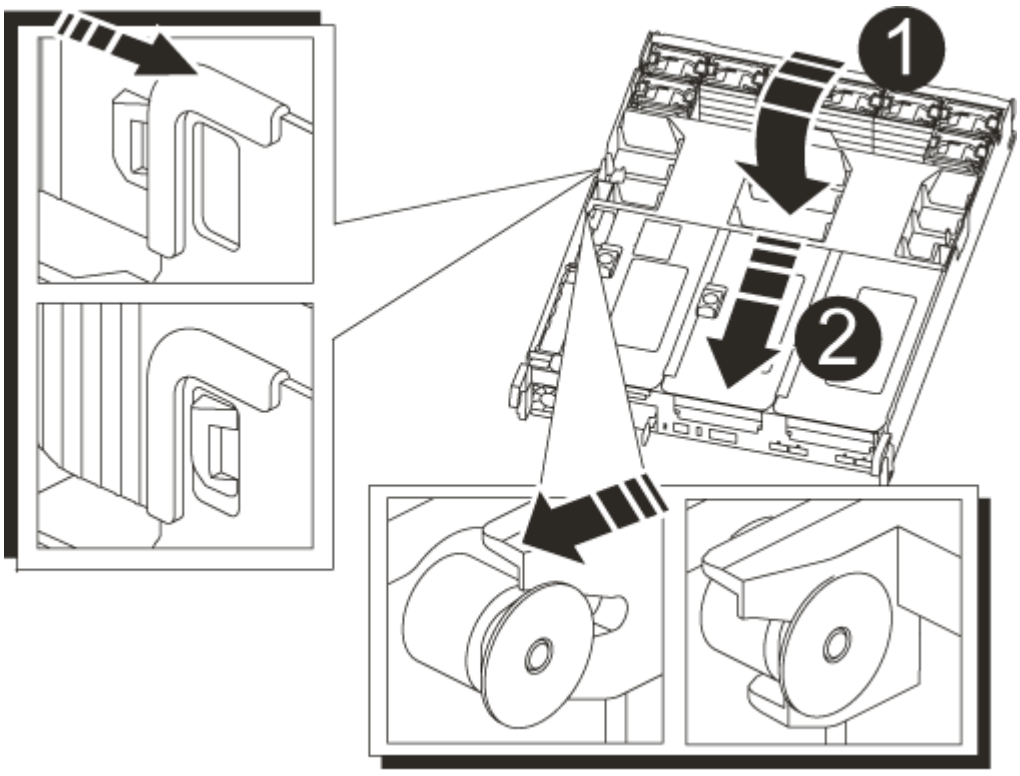

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - ASA A800

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be

resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

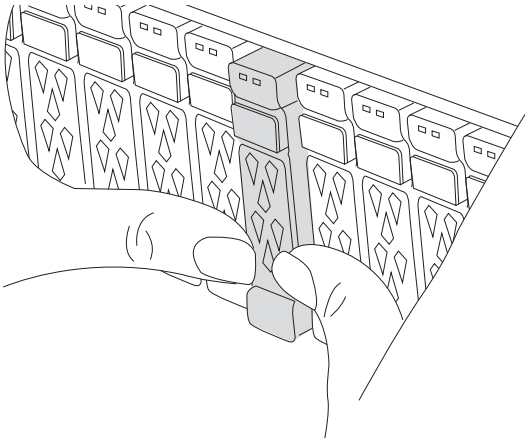

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

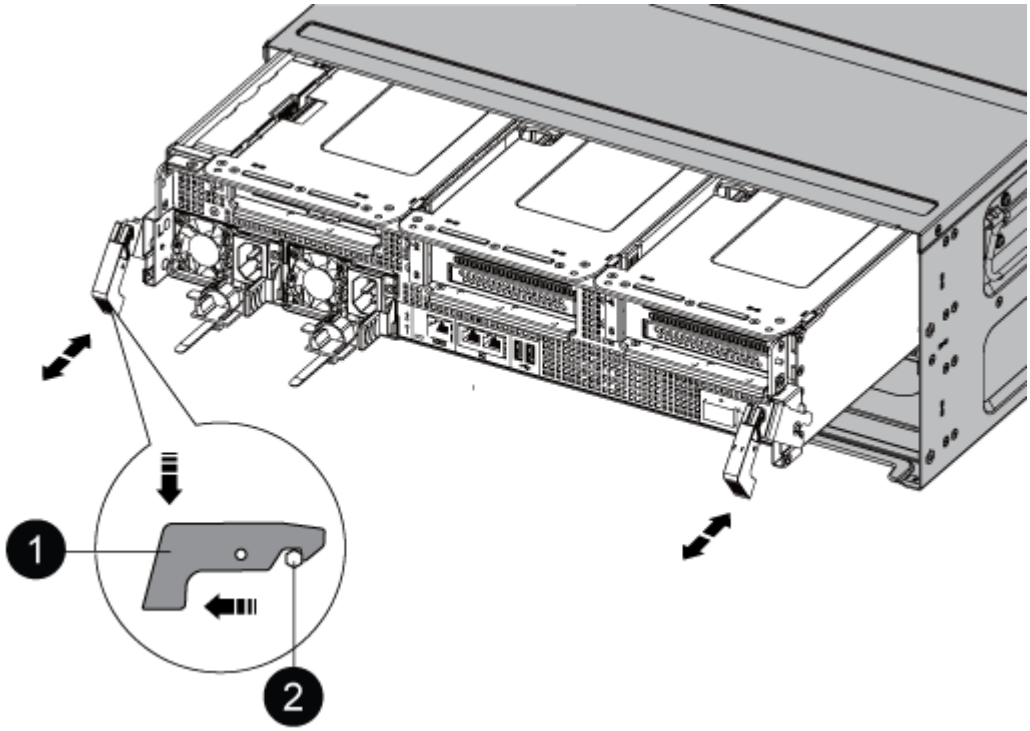

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Set the controller module aside in a safe place.

Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module.

1. Open the air duct cover and locate the NVDIMM battery in the riser.

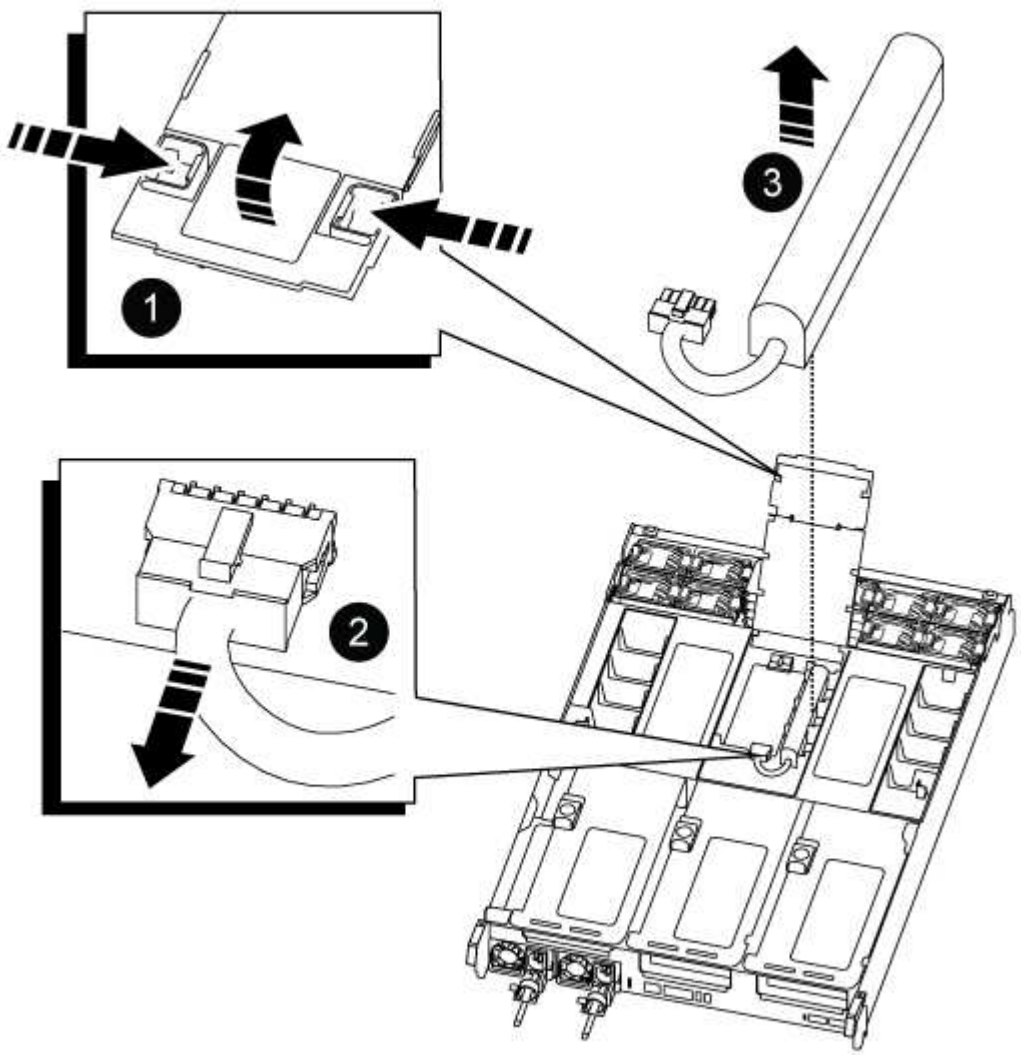

1	Air duct riser

2	NVDIMM battery plug
3	NVDIMM battery pack

Attention: The NVDIMM battery control board LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

1. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
2. Grasp the battery and lift the battery out of the air duct and controller module, and then set it aside.
3. Remove the replacement battery from its package.
4. Install the replacement battery pack in the NVDIMM air duct:
   a. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   b. Plug the battery plug into the riser socket and make sure that the plug locks into place.
5. Close the NVDIMM air duct.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace PCIe card - ASA A800

To replace a PCIe card, you must disconnect the cables from the cards, remove the SFP and QSFP modules from the cards before removing the riser, reinstall the riser, and then reinstall the SFP and QSFP modules before cabling the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

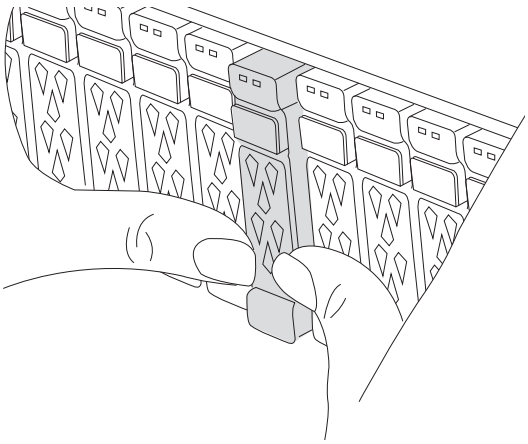

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

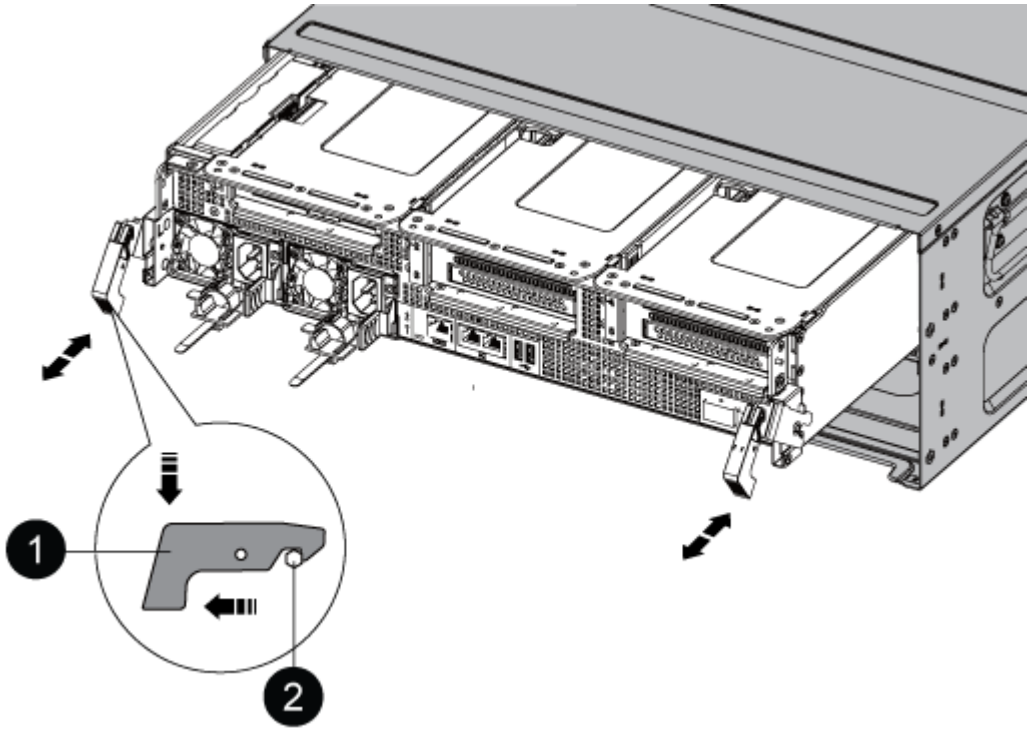

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

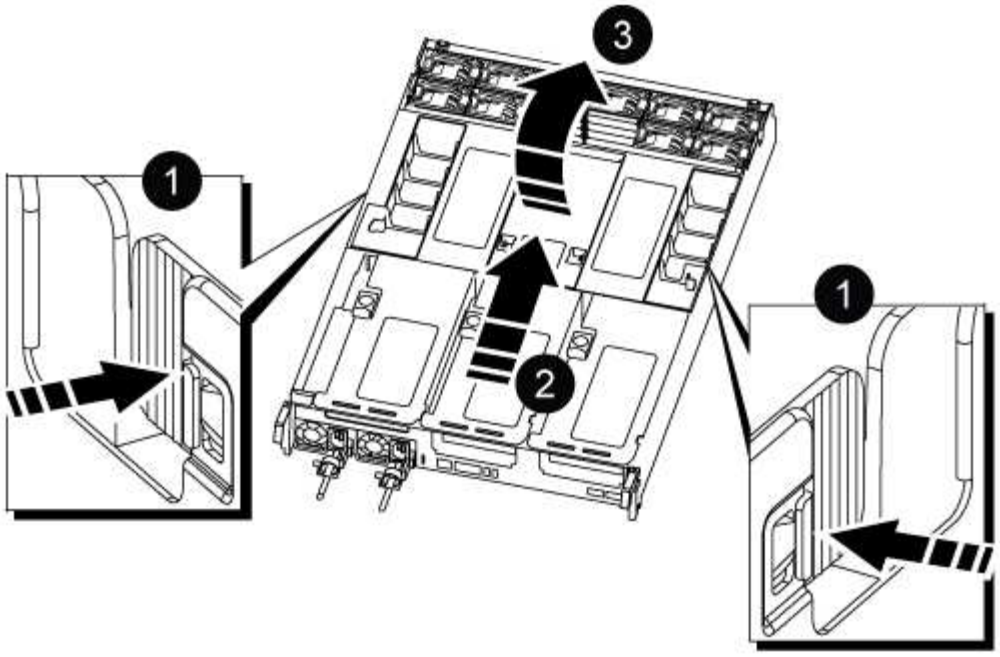

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace a PCIe card

To replace a PCIe card, you must remove the cabling and any QSFPs and SFPs from the ports on the PCIe cards in the target riser, remove the riser from the controller module, remove and replace the PCIe card, reinstall the riser and any QSFPs and SFPs onto the ports, and cable the ports.

1. Determine if the card you are replacing is from Riser 1 or if it is from Riser 2 or 3.
   - If you are replacing the 100GbE PCIe card in Riser 1, use Steps 2 - 3 and Steps 6 - 7.
   - If you are replacing a PCIe card from Riser 2 or 3, use Steps 4 through 7.
2. Remove Riser 1 from the controller module:
   a. Remove the QSFP modules that might be in the PCIe card.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

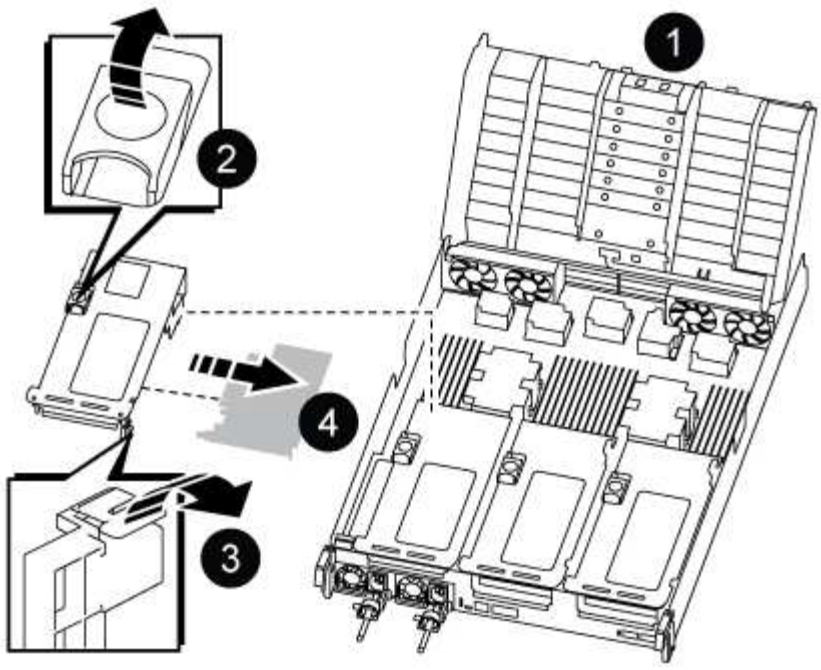

1	Air duct
2	Riser locking latch
3	Card locking bracket
4	Riser 1 (left riser) with 100GbE PCIe card in slot 1.

3. Remove the PCIe card from Riser 1:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. Remove the PCIe card from the riser.
4. Remove the PCIe riser from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

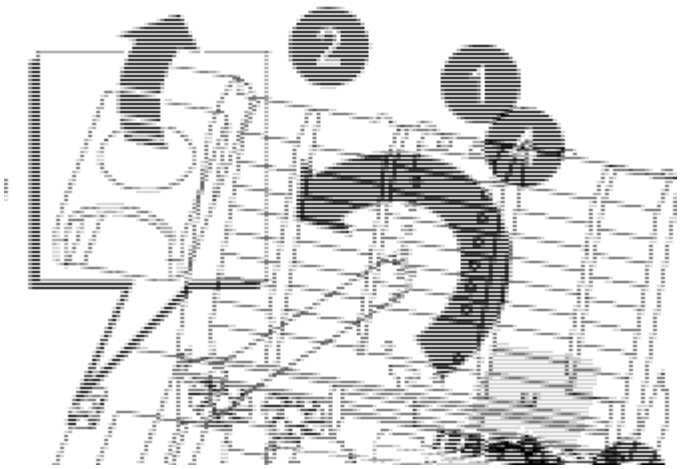

1	Air duct
2	Riser 2 (middle riser) or 3 (right riser) locking latch
3	Card locking bracket
4	Side panel on riser 2 or 3
5	PCIe cards in riser 2 or 3

5. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe cards.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. Swing the side panel off the riser.
   d. Remove the PCIe card from the riser.
6. Install the PCIe card into the same slot in the riser:
   a. Align the card with the card socket in the riser, and then slide it squarely into the socket in the riser.

      Make sure that the card is completely and squarely seated into the riser socket.

   b. For Riser 2 or 3, close the side panel.
   c. Swing the locking latch into place until it clicks into the locked position.
7. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the

controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

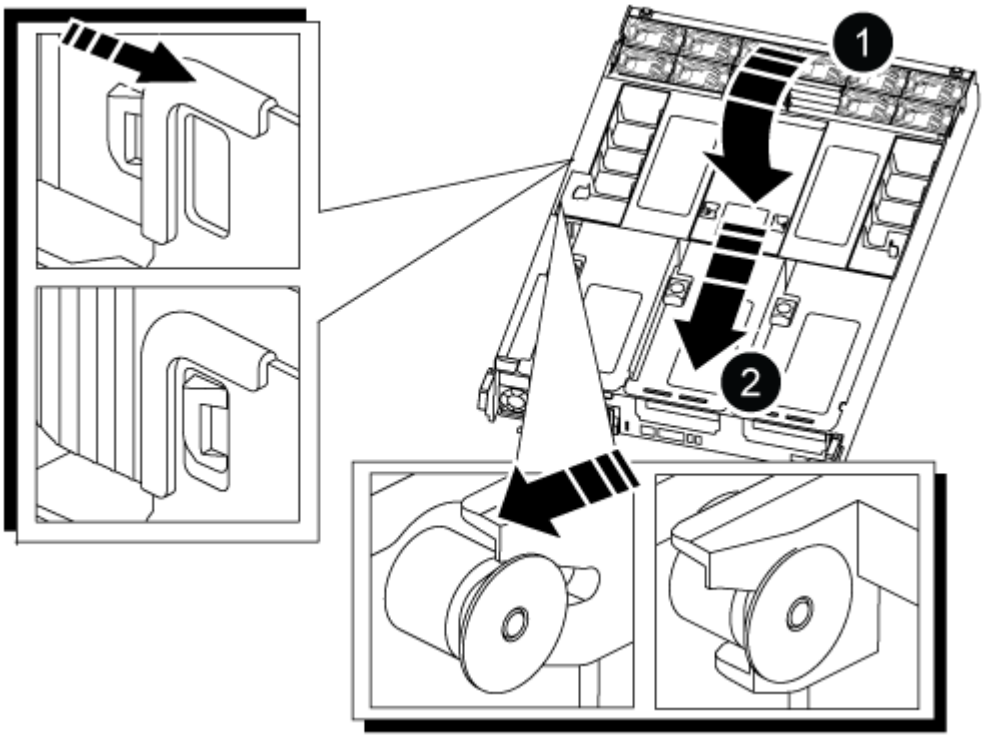

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
5. Plug the system cables and transceiver modules into the controller module and reinstall the cable management device.
6. Plug the power cables into the power supplies and reinstall the power cable retainers.

   The controller module begins to boot as soon as it is connected to power. Be prepared to interrupt the boot process.

   > If your system has DC power supplies, make sure the thumbscrews on the power supply cable are tight.

7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace power supply - ASA A800

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

About this task

This procedure is written for replacing one PSU at a time.

> It is a best practice to replace the PSU within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded PSU until the PSU is replaced.

> Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
   b. Unplug the power cable from the power source.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

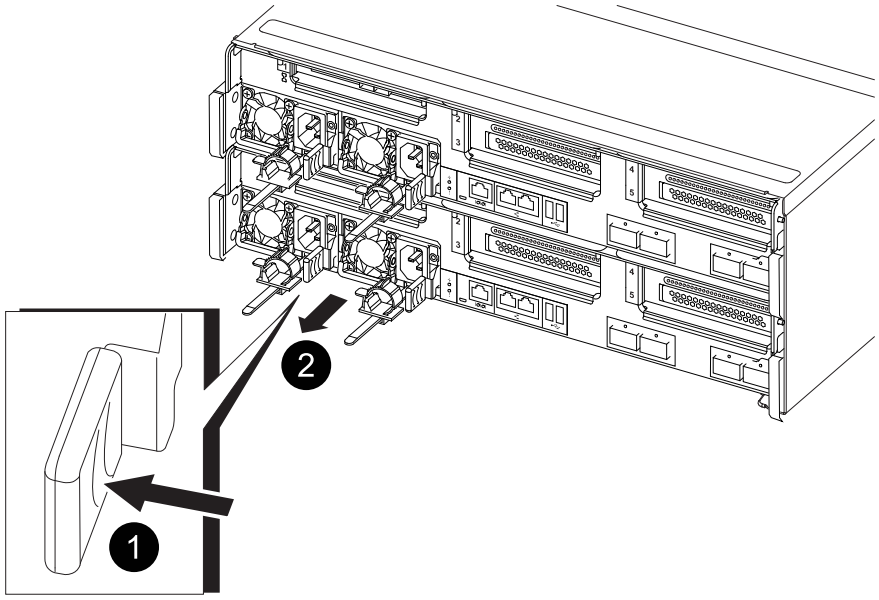

1	Blue PSU locking tab
2	Power supply

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU and the power source.
   b. Secure the power cable to the PSU using the power cable retainer.

   Once power is restored to the PSU, the status LED should be green.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

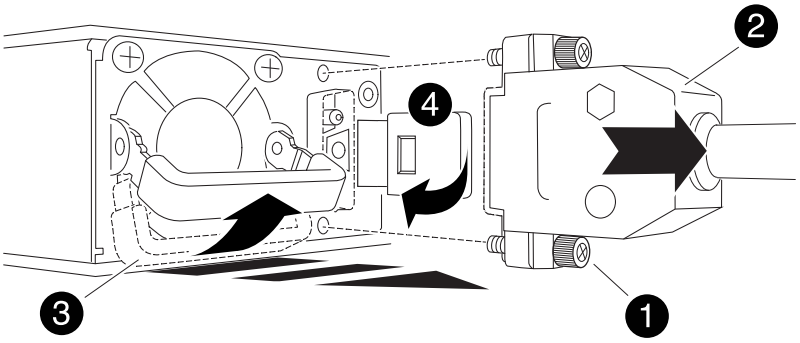

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle

4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA A800

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

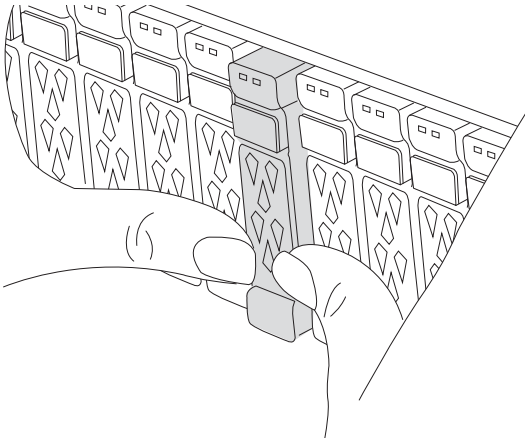

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

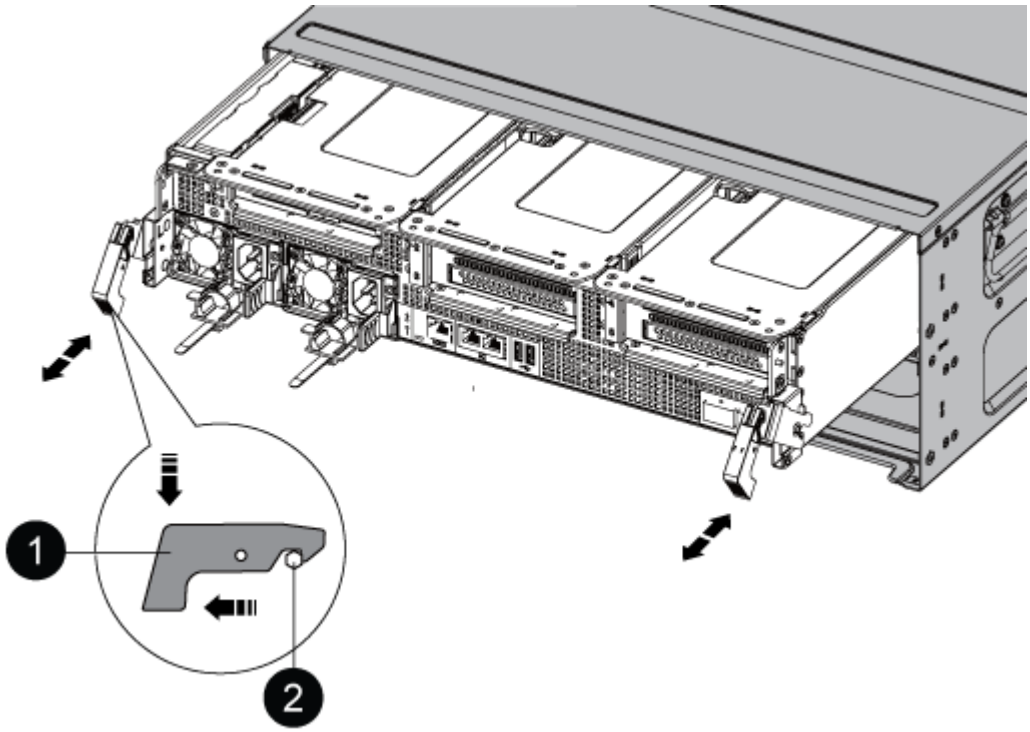

1	Locking latch
2	Locking pin

1. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

2. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

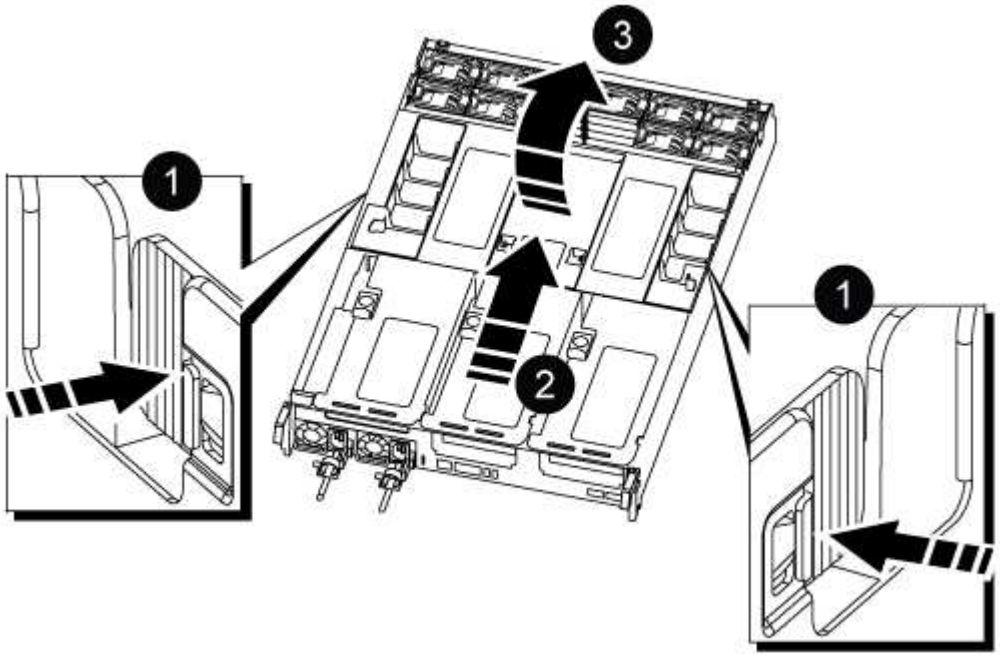

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace the RTC battery

Original controller

1. Remove PCIe riser 2 (middle riser) from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

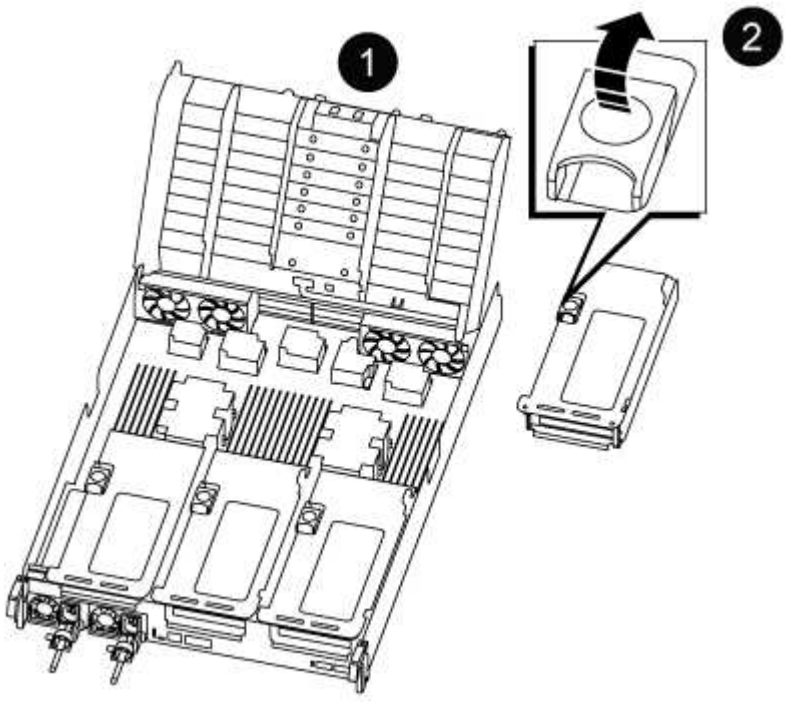

1	Air duct
2	Riser 2 (middle riser) locking latch

2. Locate the RTC battery under Riser 2.

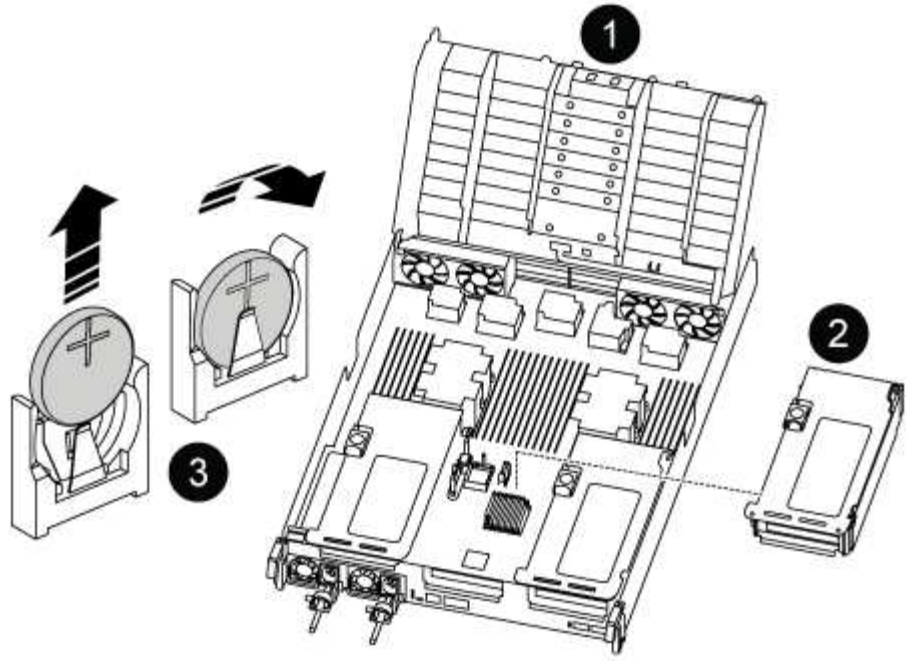

1	Air duct
2	Riser 2
3	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
7. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

VER2 controller

1. Locate the RTC battery near the DIMMs.

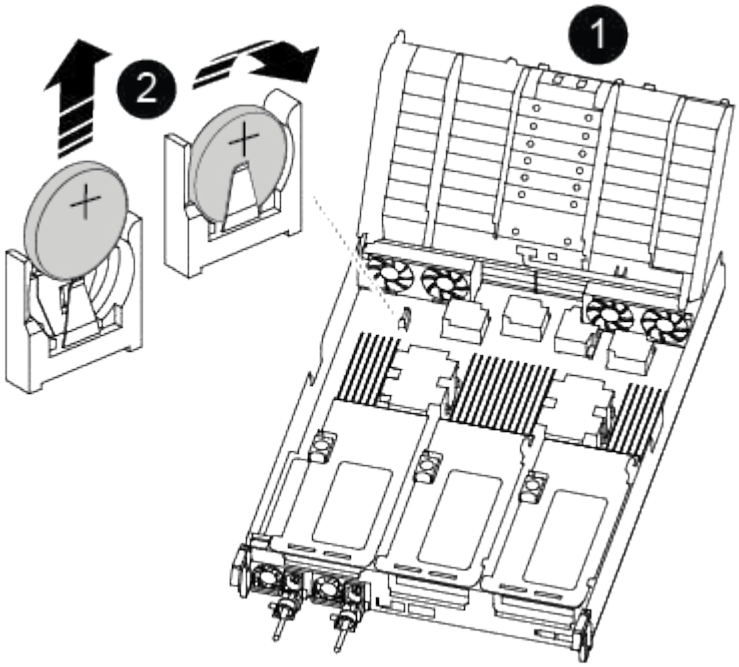

1	Air duct
2	RTC battery and housing

2. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

3. Remove the replacement battery from the antistatic shipping bag.
4. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
5. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Halt the controller at the LOADER prompt.
5. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA A900 systems

Install and setup

Start here: Choose your installation and setup experience

You can choose from different content formats to guide you through installing and setting up your new storage system.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

Quick steps - ASA A900

The quick guide provides graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this content if you are familiar with installing NetApp systems.

Use the xref:./asa900/AFF A900 Installation and Setup Instructions

The ASA A900 uses the same installation procedure as the AFF A900 system.

Video steps - ASA A900

The following video shows how to install and cable your new system.

Animation - AFF A900 Installation and setup instructions

The ASA A900 uses the same installation procedure as the AFF A900 system.

Detailed steps - ASA 900

This page provides detailed step-by-step instructions for installing a typical NetApp system. Use this article if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your system, you need to create an account on the NetApp Support Site, register your system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the NetApp Hardware Universe for information about site requirements as well as additional information on your configured system.

What you need

You might also want to have access to the ONTAP 9 Release Notes for your version of ONTAP for more information about this system.

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
25 GbE data Cable	X66240A-05 (112-00639), 0.5m X66240A-2 (112-00598), 2m X66240A-5 (112-00600), 5m		Network cable
32 Gb FC (SFP+ Op)	X66250-2 (112-00342), 2m X66250-5 (112-00344), 5m X66250-15 (112-00346), 15m		FC optical network cable
40 GbE network cable	X66100-1 (112-00542), 1m X66100-3 (112-00543), 3m X66100-5 (112-00544), 5m		Ethernet data, cluster network

Type of cable…	Part number and length	Connector type	For…
100 GbE cable	X66211B-1 (112-00573), 1m X66211B-2 (112-00574), 2m X66211B-5 (112-00576), 5m		Network, NVME storage, Ethernet data, cluster network
Optical cables	X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		FC optical network
Cat 6, RJ-45 (order dependent)	Part numbers X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network and Ethernet data
Micro-USB console cable	Not applicable		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	Not applicable		Powering up the system

4. Review the ONTAP Configuration Guide and collect the required information listed in that guide.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

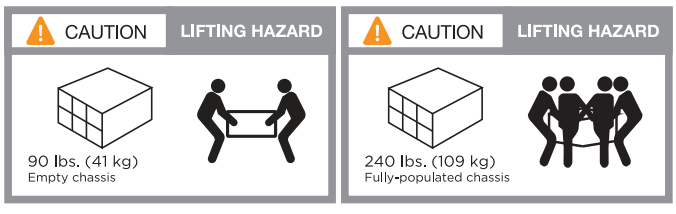

3. Attach cable management devices (as shown).

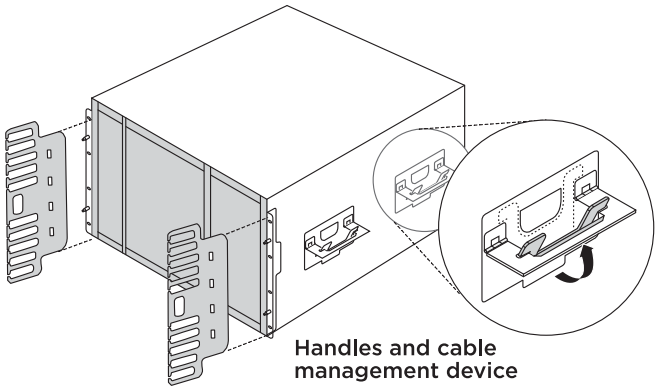

4. Place the bezel on the front of the system.

The following diagram shows a representation of what a typical system looks like and where the major components are located at the rear of the system:

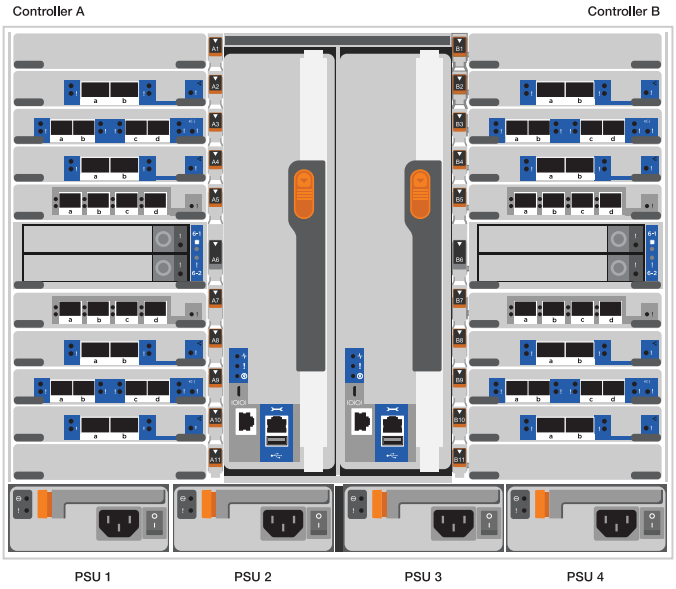

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Two-node switchless cluster

Management network, data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

Before you begin

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all networking module ports.

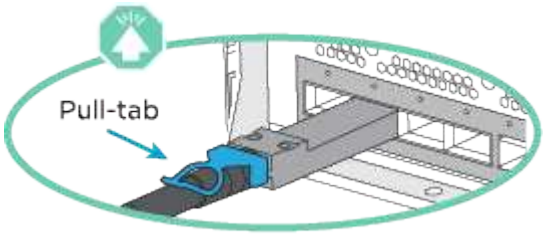

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Cable a two-node switchless cluster

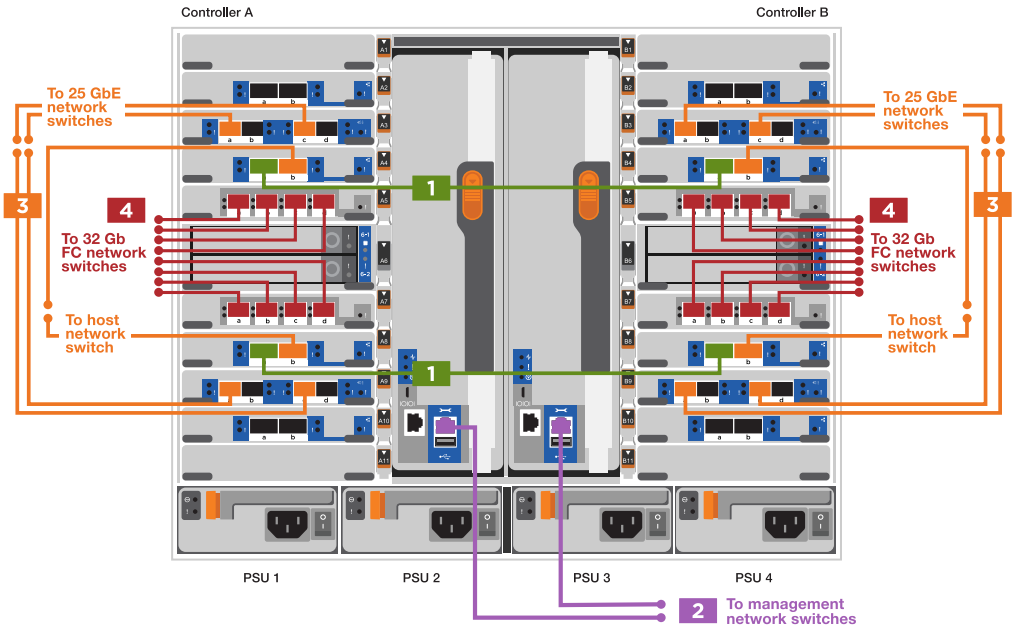

Step	Perform on each controller
1	Cable cluster interconnect ports: • Slot A4 and B4 (e4a) • Slot A8 and B8 (e8a) 
2	Cable controller management (wrench) ports. 
3	Cable 25 GbE network switches: Ports in slot A3 and B3 (e3a and e3c) and slot A9 and B9 (e9a and e9c) to the 25 GbE network switches.  40GbE host network switches: Cable host-side b ports in slot A4 and B4 (e4b) and slot A8 and B8 (e8b) to the host switch. 
4	Cable 32 Gb FC connections: Cable ports in slot A5 and B5 (5a, 5b, 5c, and 5d) and slot A7 and B7 (7a, 7b, 7c, and 7d) to the 32 Gb FC network switches. 

Step	Perform on each controller
5	• Strap the cables to the cable management arms (not shown). • Connect the power cables to the PSUs and connect them to different power sources (not shown). PSU 1 and 3 provide power to all side A components, while PSU2 and PSU4 provide power to all side B components. 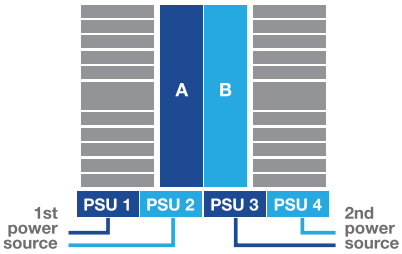

Option 2: Switched cluster

Management network, data network, and management ports on the controllers are connected to switches. The cluster interconnect and HA ports are cabled on to the cluster/HA switch.

Before you begin

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all networking module ports.

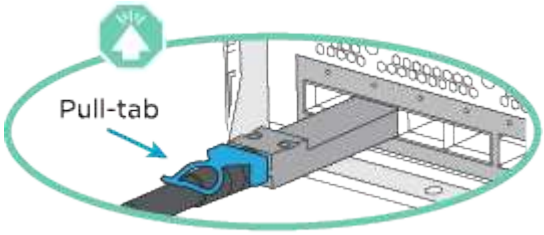

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it over and try again.

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Cable a switched cluster

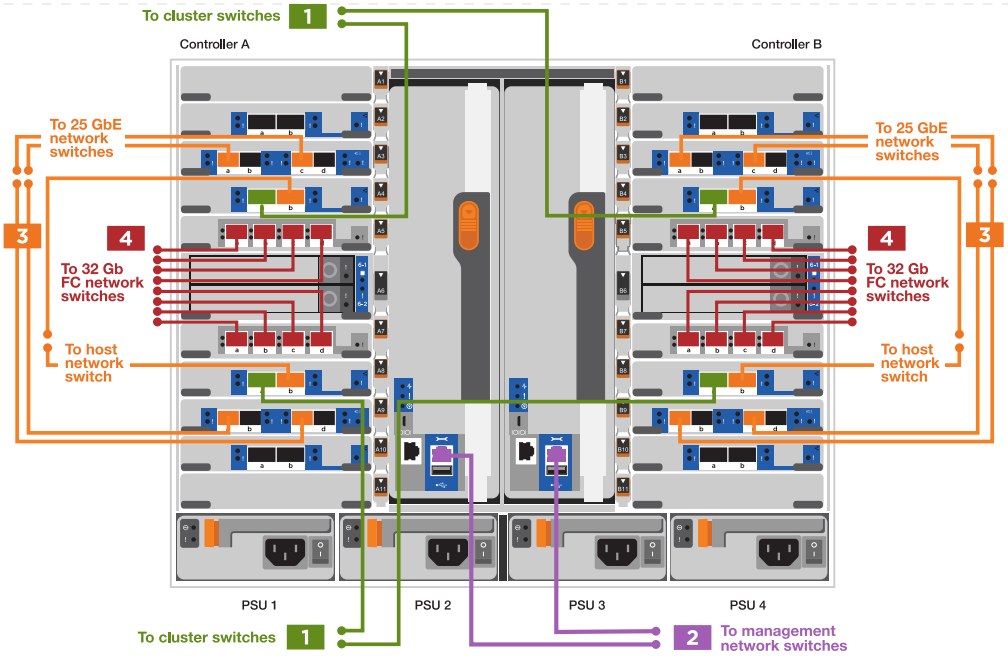

Step	Perform on each controller
1	Cable cluster interconnect a ports: • Slot A4 and B4 (e4a) to the cluster network switch. • Slot A8 and B8 (e8a) to the cluster network switch. 
2	Cable controller management (wrench) ports. 

Step	Perform on each controller
3	Cable 25GbE network switches: Ports in slot A3 and B3 (e3a and e3c) and slot A9 and B9 (e9a and e9c) to the 25 GbE network switches.  40GbE host network switches: Cable host-side b ports in slot A4 and B4 (e4b) and slot A8 and B8 (e8b) to the host switch. 
4	Cable 32 Gb FC connections: Cable ports in slot A5 and B5 (5a, 5b, 5c, and 5d) and slot A7 and B7 (7a, 7b, 7c, and 7d) to the 32 Gb FC network switches. 
5	• Strap the cables to the cable management arms (not shown). • Connect the power cables to the PSUs and connect them to different power sources (not shown). PSU 1 and 3 provide power to all side A components, while PSU2 and PSU4 provide power to all side B components.  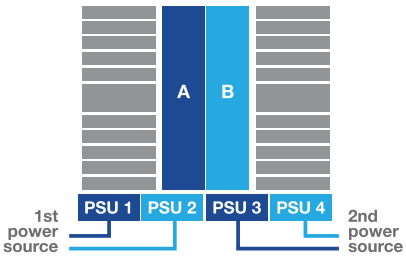

Step 4: Cable controllers to drive shelves

Cable either a single NS224 drive shelf or two NS224 drive shelves to your controllers.

Option 1: Cable the controllers to a single NS224 drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

Before you begin

- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the storage modules are up, while the pull tabs on the shelves are down.

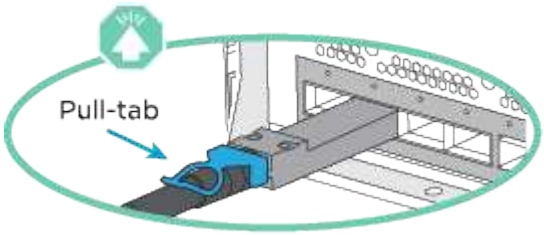

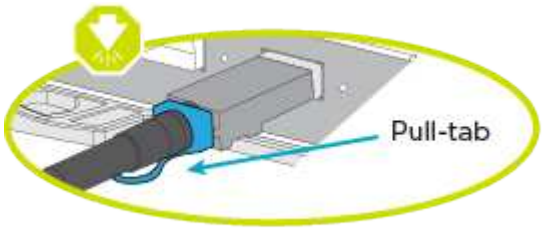

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

1. Use the following animation or drawings to cable your controllers to a single NS224 drive shelf.

   Animation - Cable a single NS224 shelf

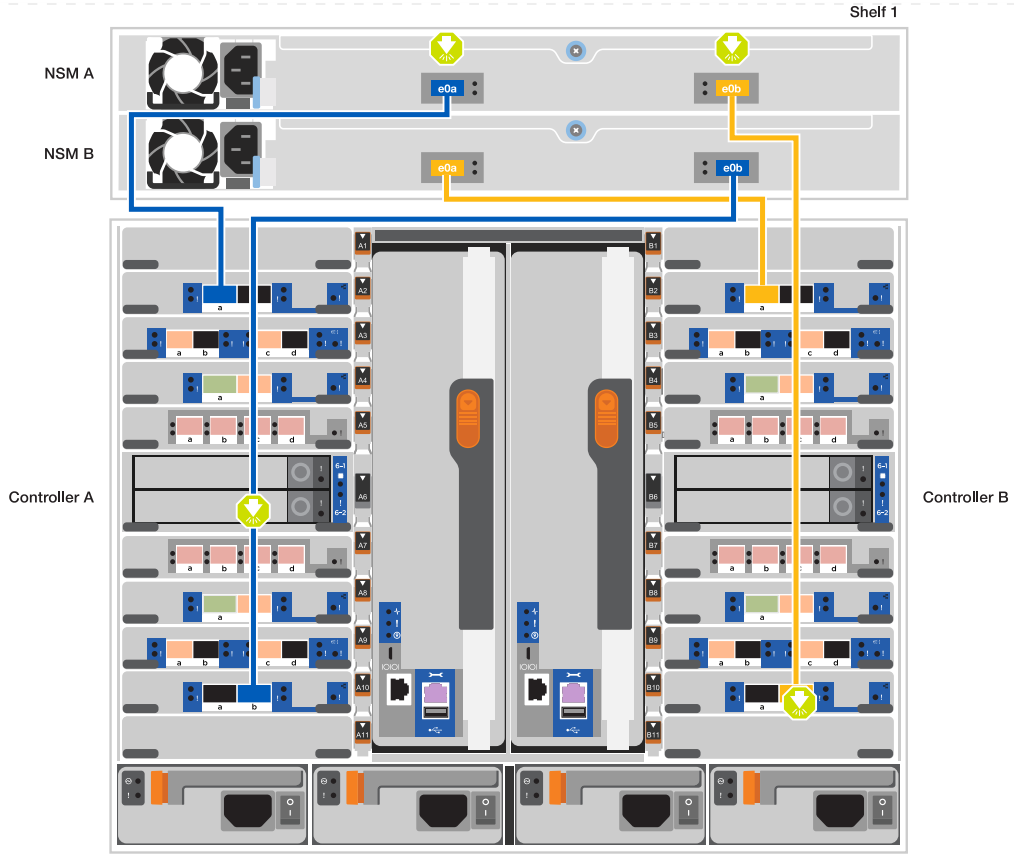

Step	Perform on each controller
1	• Connect controller A port e2a to port e0a on NSM A on the shelf. • Connect controller A port e10b to port e0b on NSM B on the shelf. 100 GbE cable
2	• Connect controller B port e2a to port e0a on NSM B on the shelf. • Connect controller B port e10b to port e0b on NSM A on the shelf. 100 GbE cable

Option 2: Cable the controllers to two NS224 drive shelves

You must cable each controller to the NSM modules on the NS224 drive shelves.

Before you begin

- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the storage modules are up, while the pull tabs on the shelves are down.

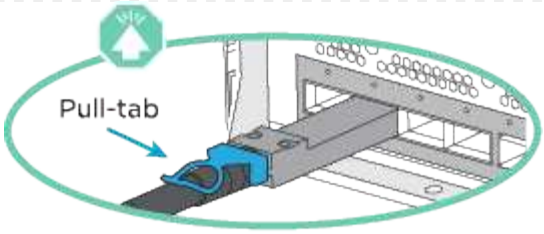

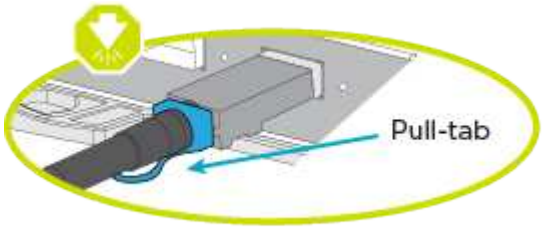

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

1. Use the following animation or diagram to cable your controllers to two NS224 drive shelves.

Animation - Cable two NS224 shelves

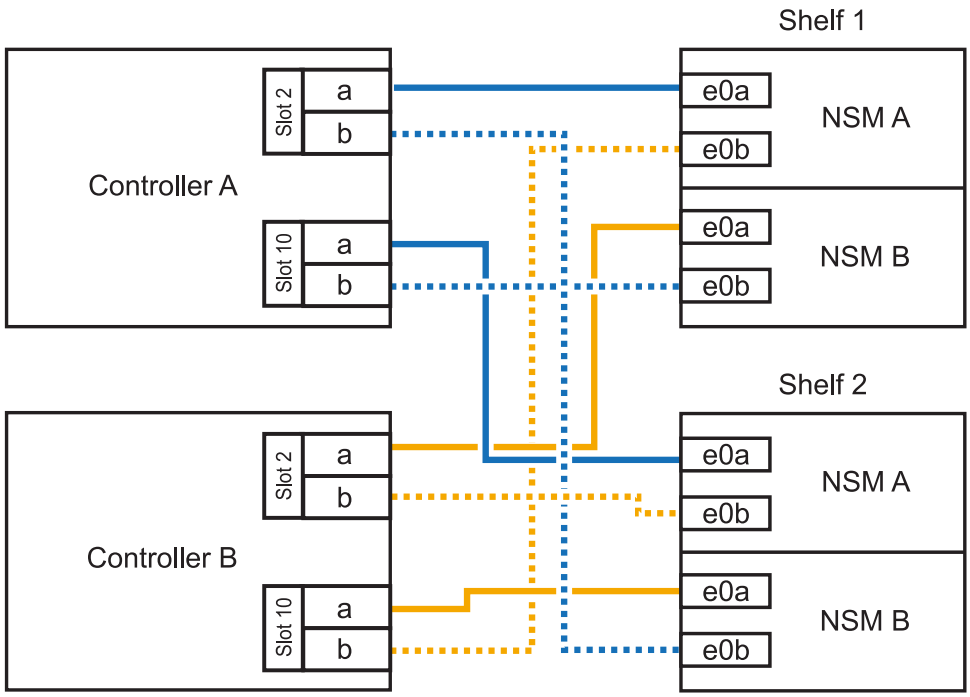

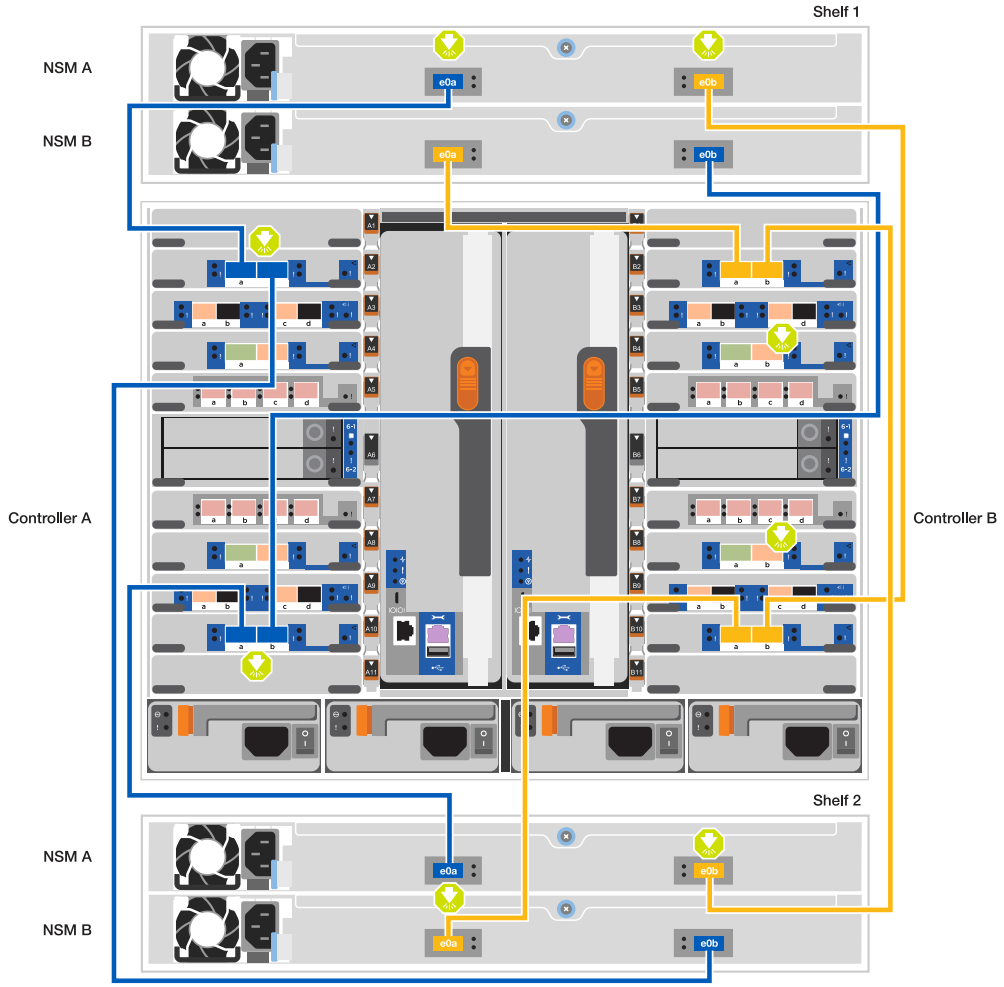

Step	Perform on each controller
1	• Connect controller A port e2a to NSM A e0a on shelf 1. • Connect controller A port e10b to NSM B e0b on shelf 1. • Connect controller A port e2b to NSM B e0b on shelf 2. • Connect controller A port e10a to NSM A e0a on shelf 2.  100 GbE cable
2	• Connect controller B port e2a to NSM B e0a on shelf 1. • Connect controller B port e10b to NSM A e0b on shelf 1. • Connect controller B port e2b to NSM A e0b on shelf 2. • Connect controller B port e10a to NSM B e0a on shelf 2.  100 GbE cable

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: If network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

1. Use the following animation or drawing to set one or more drive shelf IDs:

   The NS224 shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located. see Change a shelf ID - NS224 shelves for detailed instructions.

   Animation - Set NVMe drive shelf IDs

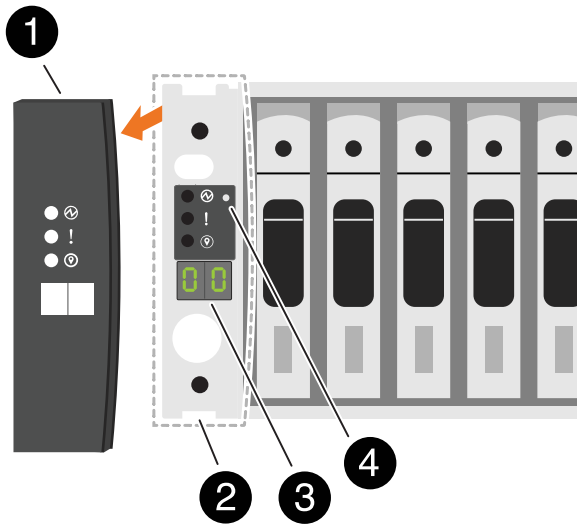

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID LED
4	Shelf ID setting button

2. Turn on the power switches on the power supplies to both nodes.

   Animation - Turn on the power to the controllers

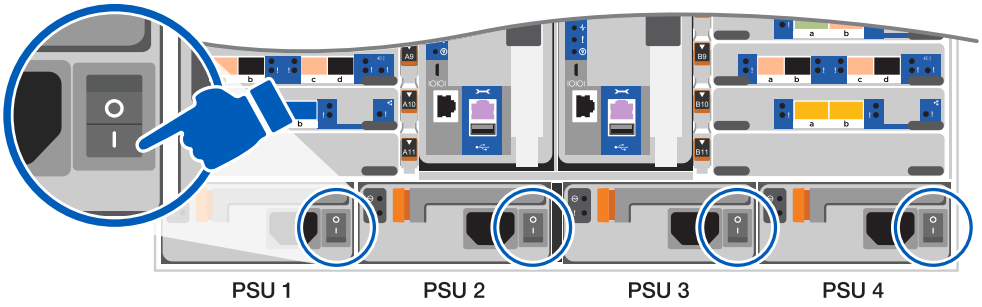

Initial booting may take up to eight minutes.

3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

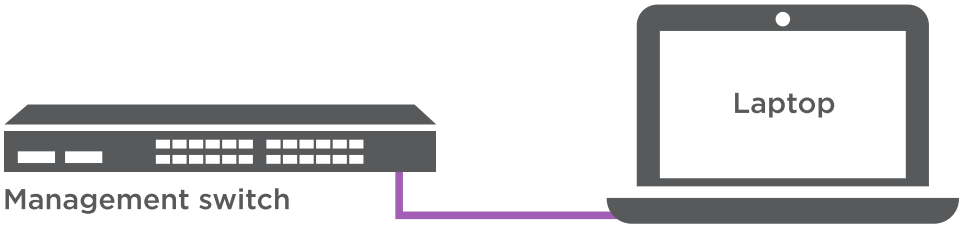

5. Select an ONTAP icon listed to discover:

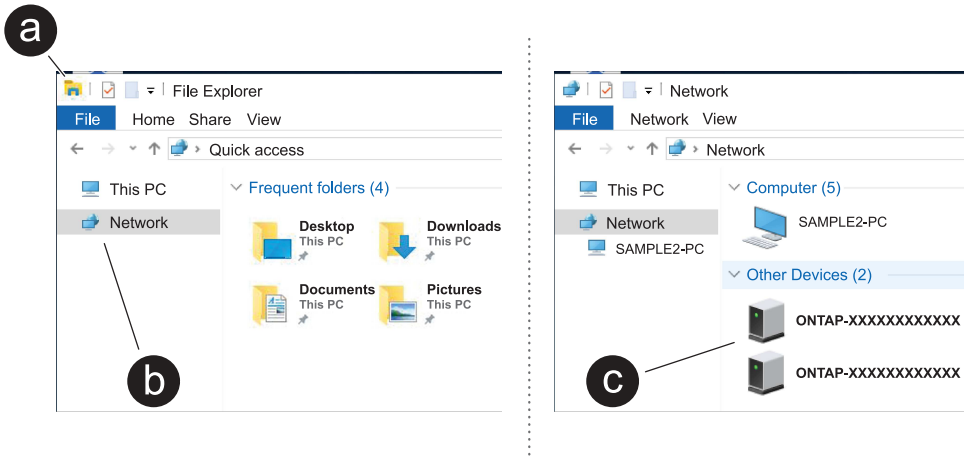

   a. Open File Explorer.

   b. Click **Network** in the left pane and right-click and select **refresh**.

   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

   System Manager opens.

6. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

If you are not using a Windows or Mac-based laptop or console or if auto discovery is not enabled, you must complete the configuration and setup using this task.

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet.

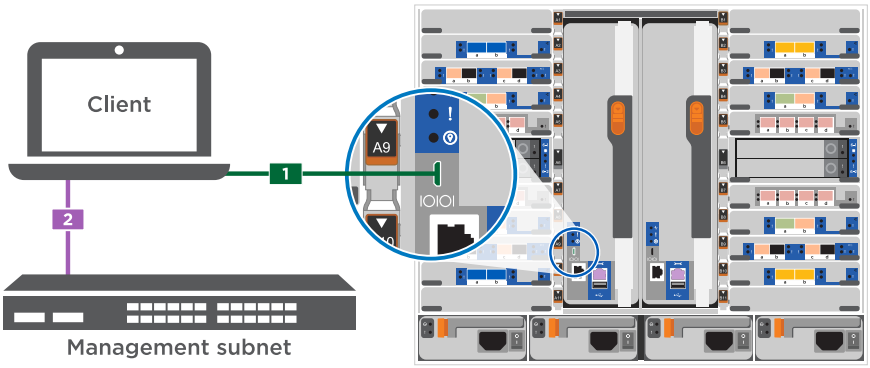

   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Use the following animation to set one or more drive shelf IDs:

   The NS224 shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located. see Change a shelf ID - NS224 shelves for detailed instructions.

   Animation - Set NVMe drive shelf IDs

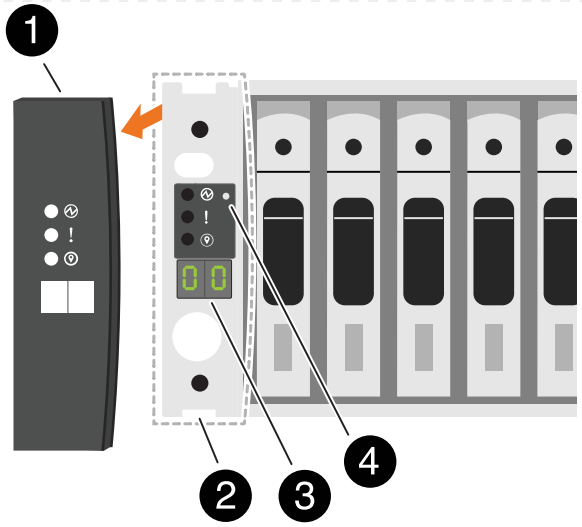

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID LED
4	Shelf ID setting button

3. Turn on the power switches on the power supplies to both nodes.

   Animation - Turn on the power to the controllers

   image:[Callout number 1] drw_a900_power-on_IEOPS-941.svg[width=500px]

   Initial booting may take up to eight minutes.

4. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.

If the management network has DHCP…	Then…
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the ONTAP Configuration Guide
6. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor
7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain ASA A900 hardware

Maintain the hardware of your ASA A900 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the ASA A900 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the ASA A900 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
DCPM	The DCPM (destage controller power module) contains the NVRAM11 battery.
Fan	The fan cools the controller.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
LED USB	The LED USB module provides connectivity to console ports and system status.
NVRAM	The NVRAM module (Non-Volatile Random Access Memory) allows the controller to retain data across power cycles or system reboots, while the NVRAM DIMM maintains NVRAM settings.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media automated recovery workflow - ASA A900

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your ASA A900 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

Review the boot media requirements

Review the requirements for boot media replacement.

Shut down the controller

Shut down the controller in your storage system when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the controller module and install the replacement boot media.

Restore the image on the boot media

Restore the ONTAP image from the partner controller.

Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - ASA A900

Before replacing the boot media in your ASA A900, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - ASA A900

Shut down the impaired controller in your ASA A900 storage system to prevent data loss and ensure system stability when replacing the boot media.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - ASA A900

The boot media in your ASA A900 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

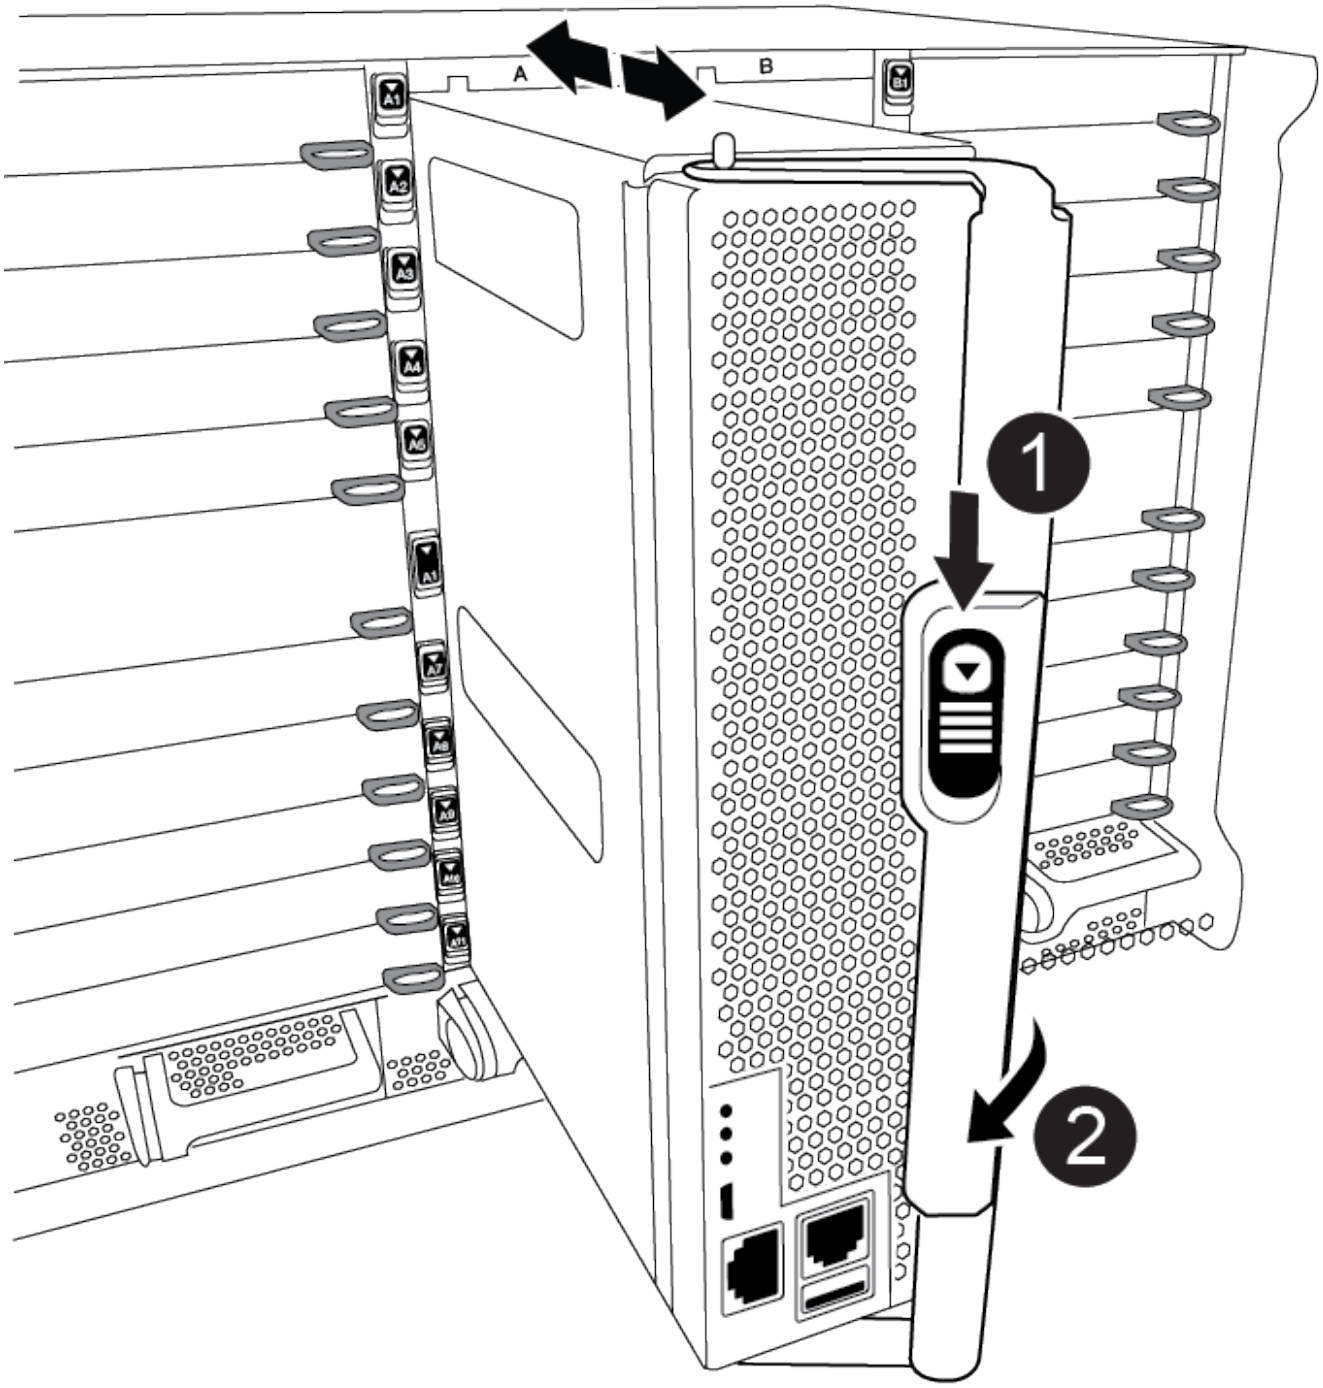

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

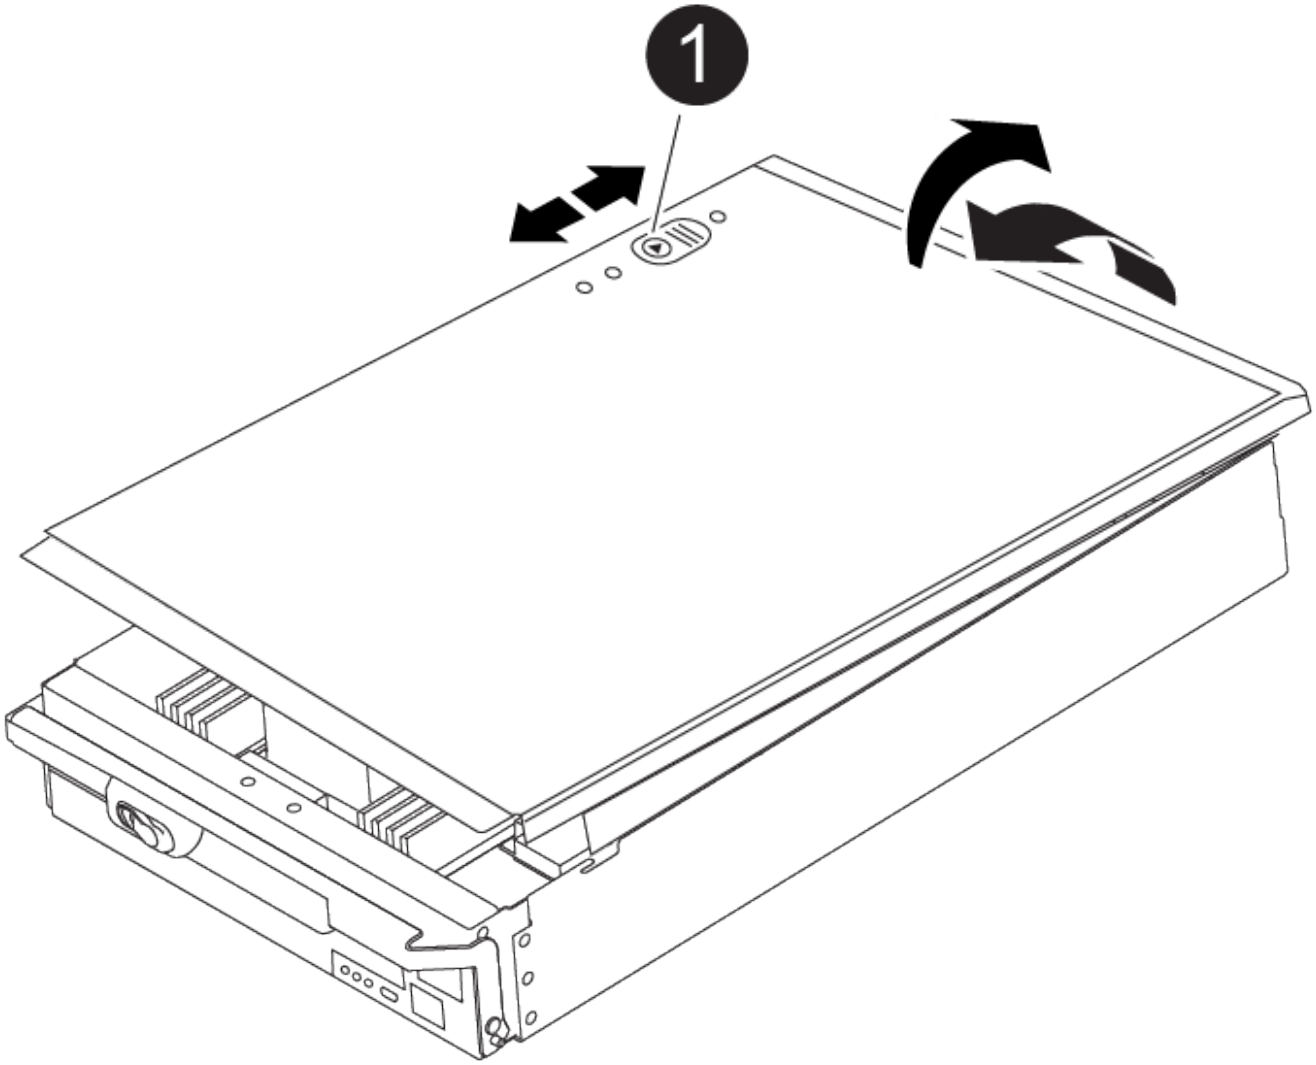

1	Controller module cover locking button

6. Replace the boot media:
   a. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

      Animation - Replace boot media

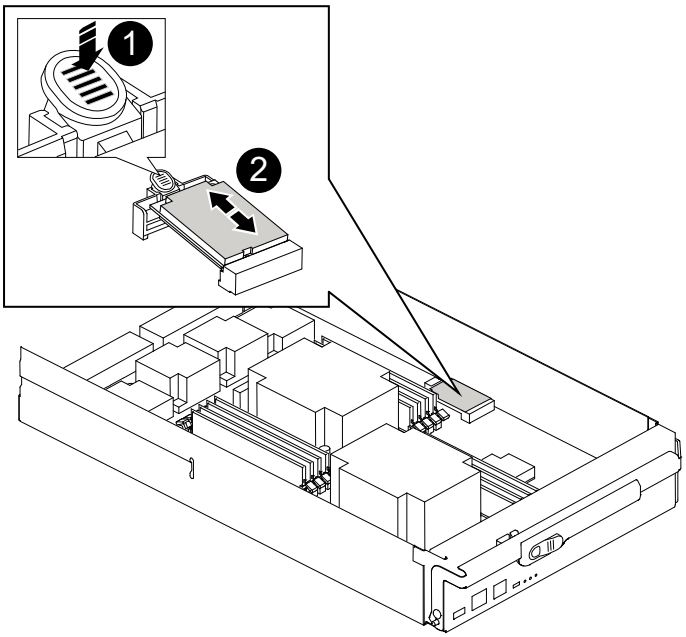

1	Press release tab
2	Boot media

b. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

c. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.

d. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

e. Push the boot media down to engage the locking button on the boot media housing.

7. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.
8. Reinstall the controller module:
   a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
   b. Recable the controller module, as needed.
   c. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, and then push the cam handle to the closed position.

      The controller begins to boot as soon as it is completely installed into the chassis.

If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

9. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t iniitator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - ASA A900

After installing the new boot media device in your ASA A900 system, you can start the automated boot media recovery process to restore the configuration from the partner node.

During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed part to NetApp - ASA A900

If a component in your ASA A900 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Replace the boot media - ASA A900

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair does not require connection to a network to restore the `var` file system. The HA pair in a single chassis has an internal e0S connection, which is used to transfer `var` config between them.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from NetApp.
- It is important that you apply the commands in these steps on the correct controller:
  - The *impaired* controller is the controller on which you are performing maintenance.
  - The *healthy* controller is the HA partner of the impaired controller.

Check encryption key support and status - ASA A900

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - ASA A900

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media and prepare for manual boot recovery - ASA A900

You must unplug the controller module, remove and open the controller module, locate and replace the boot media in the controller, and then transfer the image to the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

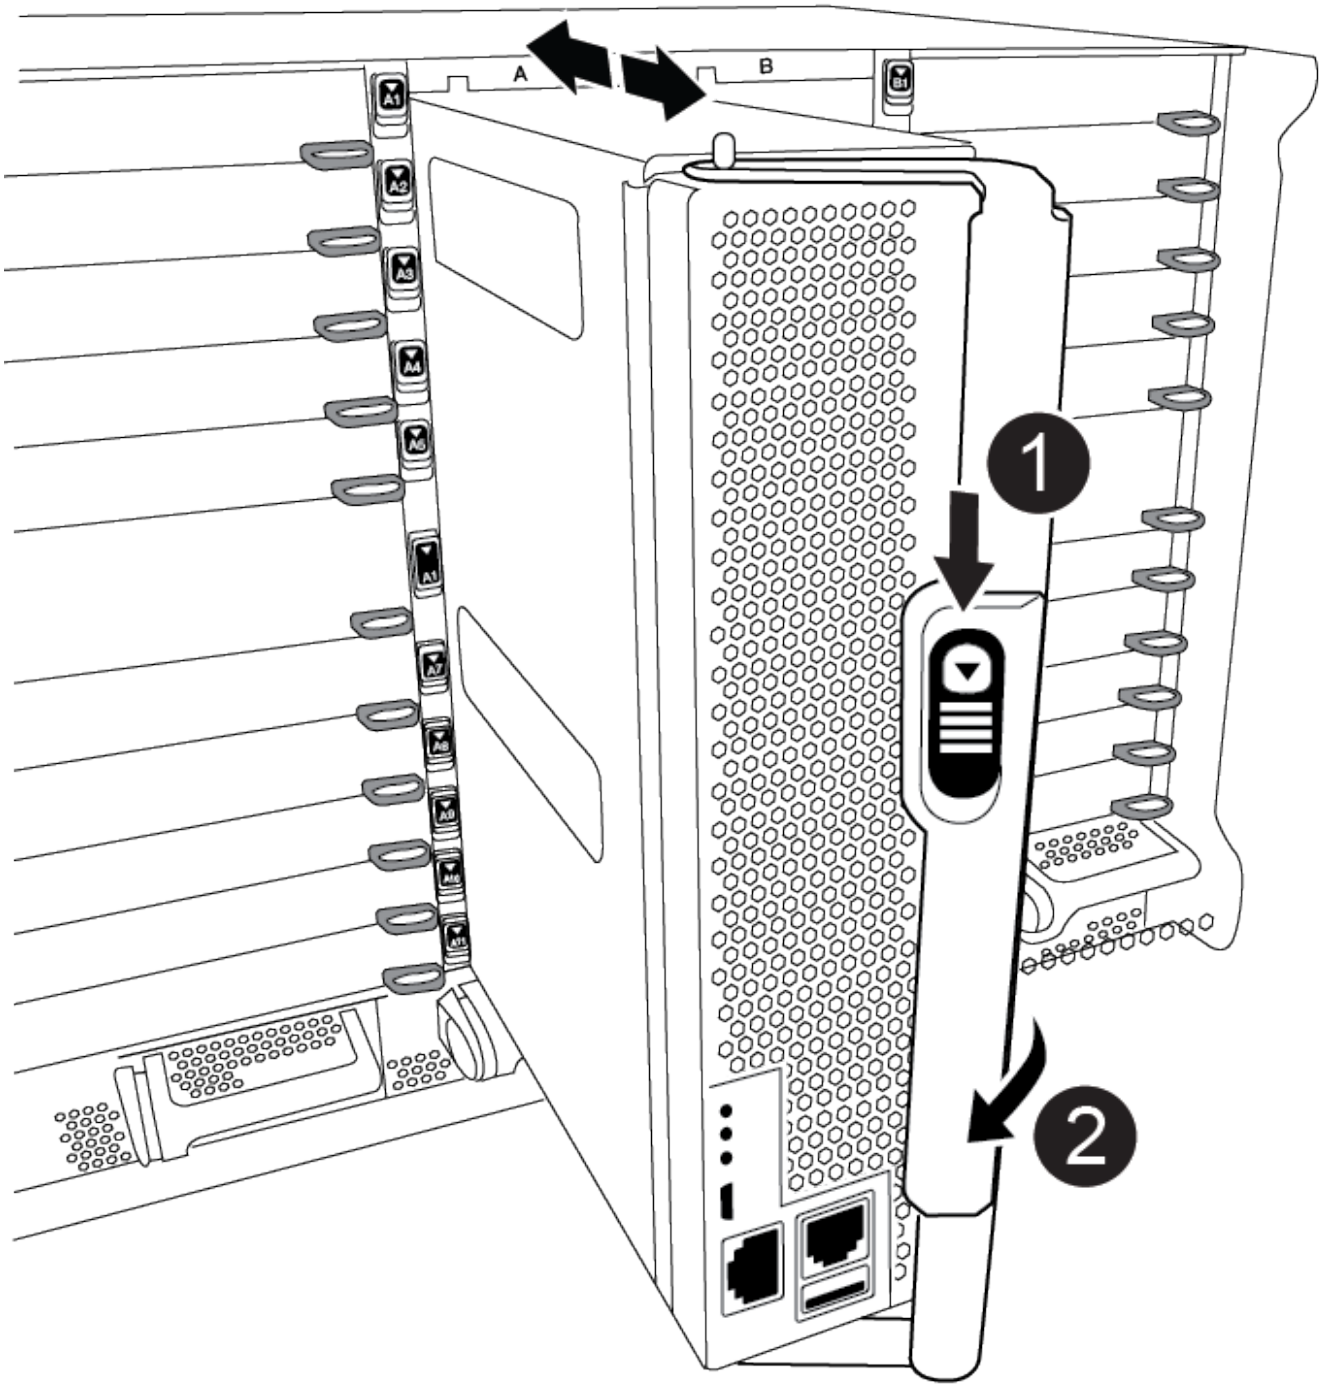

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

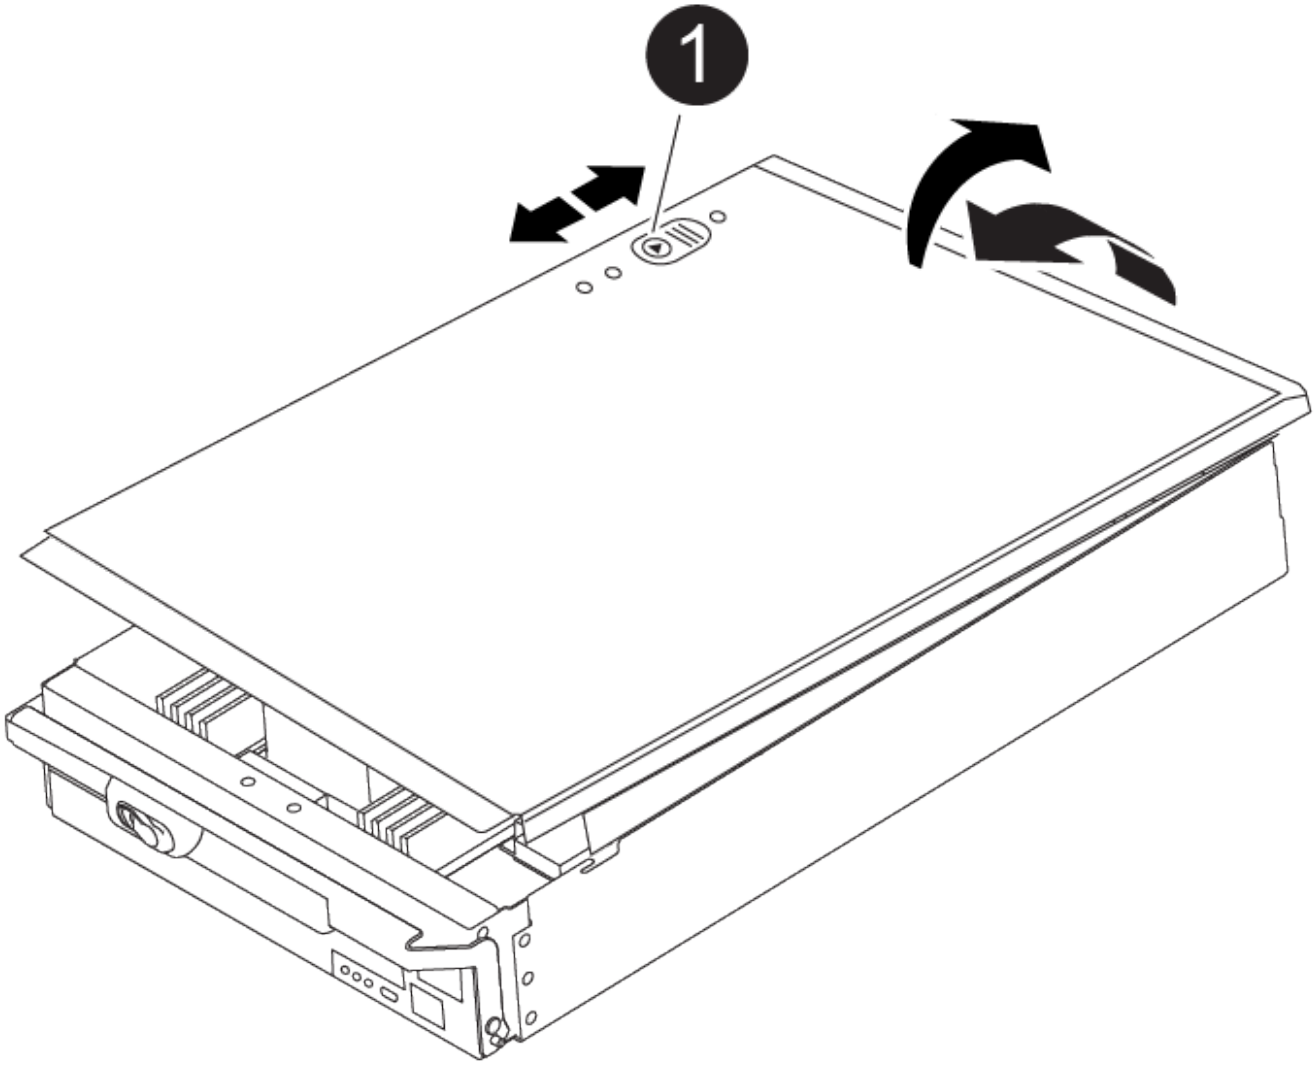

1	Controller module cover locking button

Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

   Animation - Replace boot media

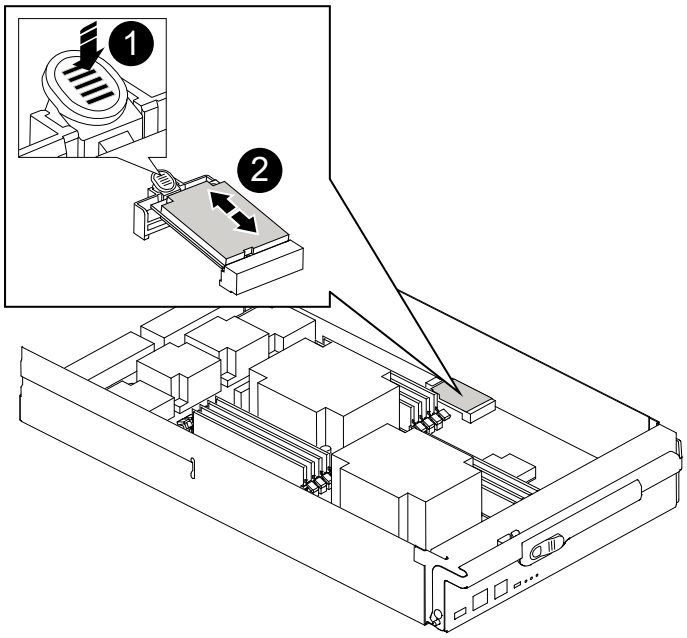

1	Press release tab
2	Boot media

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.
6. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- Download a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site. Use the

`version -v` command to display if your version of ONTAP supports NVE. If the command output displays `<10no- DARE>`, your version of ONTAP does not support NVE.

  - If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. If you have not done so, download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image from the Downloads link on the page, to your work space on your laptop.
   b. Unzip the service image.

      > If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      The USB flash drive should have the appropriate ONTAP image of what the impaired controller is running.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
3. Recable the controller module, as needed.
4. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

5. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, and then push the cam handle to the closed position.

   The controller begins to boot as soon as it is completely installed into the chassis.

6. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

7. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - ASA A900

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `...would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore encryption - ASA A900

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
**************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - ASA A900

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Replace the chassis - ASA A900

To replace the chassis, you must remove the power supplies, fans, controller modules, I/O modules, DCPM modules, and USB LED module from the impaired chassis, remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis in its place, and then install the components into the replacement chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shutdown the controllers - ASA A900

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   ```

```
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

```
Warning: Are you sure you want to halt node <node_name>? {y|n}:
```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - ASA A900

To replace the chassis, you must remove the components from the impaired chassis and install them in the replacement chassis.

Step 1: Remove the power supplies

Removing the power supplies when replacing a chassis involves turning off, disconnecting, and then removing the four power supplies from the rear of the impaired chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Press and hold the terra cotta locking button on the power supply handle, and then pull the power supply out of the chassis.

> When removing a power supply, always use two hands to support its weight.

Animation - Remove/install PSU

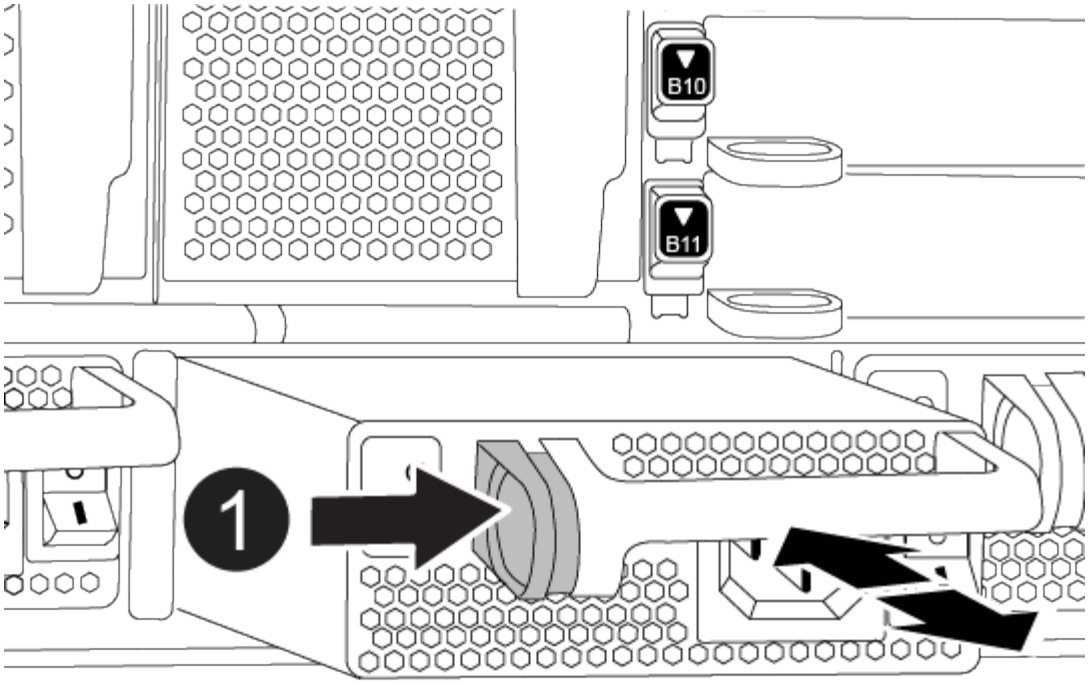

1	Locking button

4. Repeat the preceding steps for any remaining power supplies.

Step 2: Remove the fans

You must remove the six fan modules, located on in the front of the chassis, when replacing the chassis.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press the terra cotta locking button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

> The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

Animation - Remove/install fan

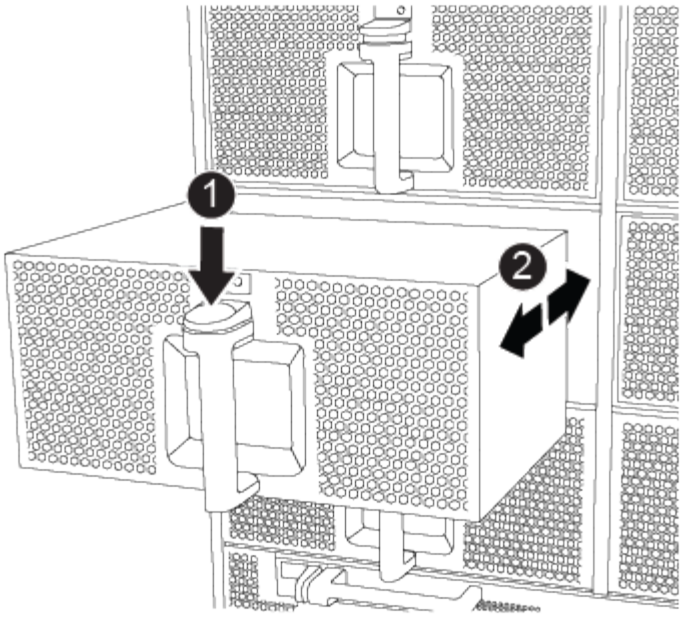

1	Terra cotta locking button
2	Slide fan in/out of chassis

4. Set the fan module aside.
5. Repeat the preceding steps for any remaining fan modules.

Step 3: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the impaired chassis.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta locking button on the cam handle downward until it unlocks.

   Animation - Remove the controller

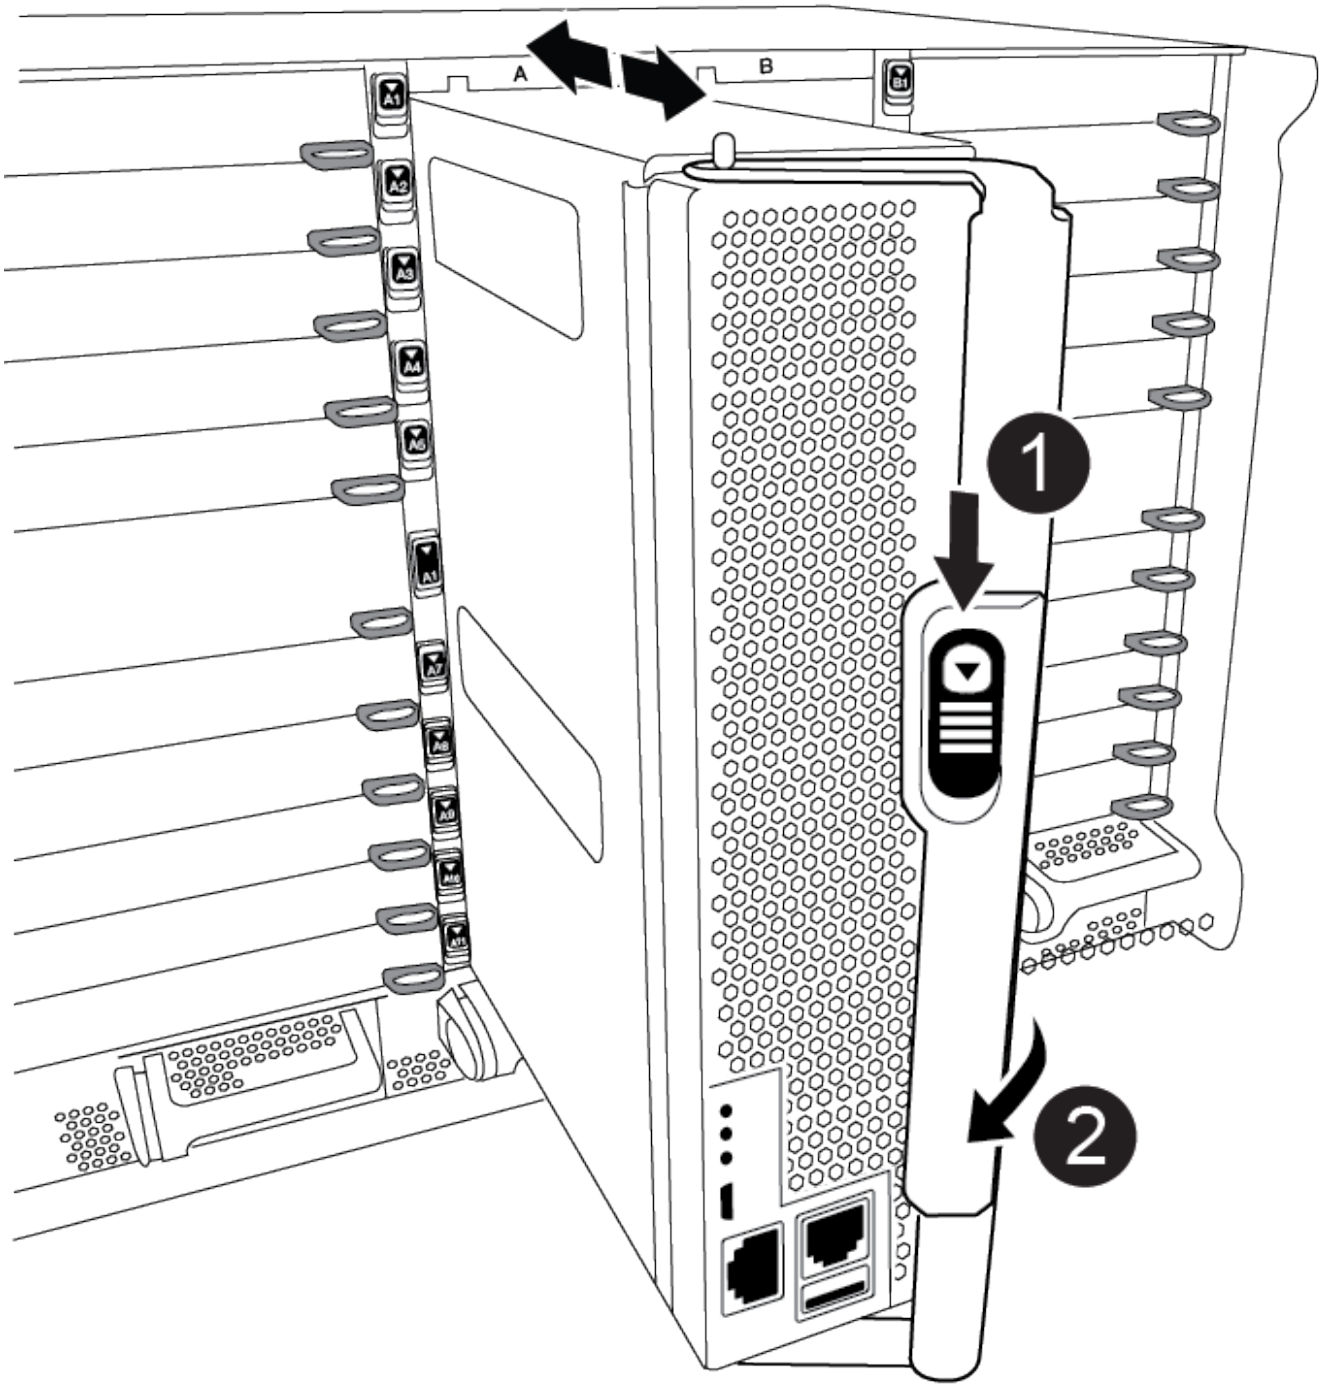

1	Cam handle locking button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Set the controller module aside in a safe place and keep track of which chassis slot it came from, so that it can be installed into the same slot in the replacement chassis..
6. Repeat these steps if you have another controller module in the chassis.

Step 4: Remove the I/O modules

To remove I/O modules from the impaired chassis, including the NVRAM modules, follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

3. Remove the target I/O module from the chassis:

   a. Depress the lettered and numbered cam locking button.

      The cam locking button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

      Animation - Remove/install I/O module

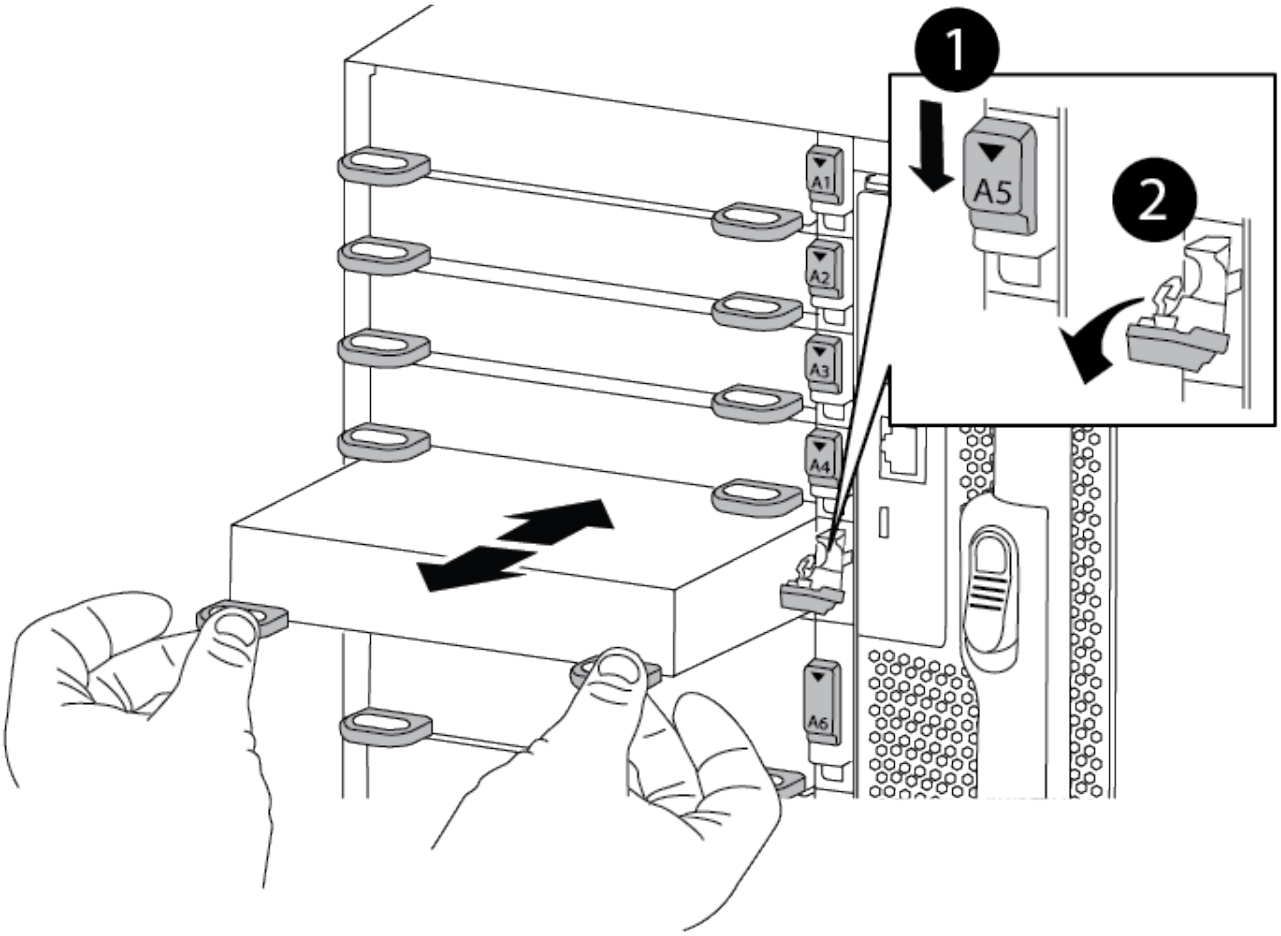

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Repeat the preceding step for the remaining I/O modules in the impaired chassis.

Step 5: Remove the de-stage controller power module

Remove the two de-stage controller power modules from the front of the impaired chassis.

1. If you are not already grounded, properly ground yourself.
2. Press the terra cotta locking button on the module handle, and then slide the DCPM out of the chassis.

   Animation - Remove/install DCPM

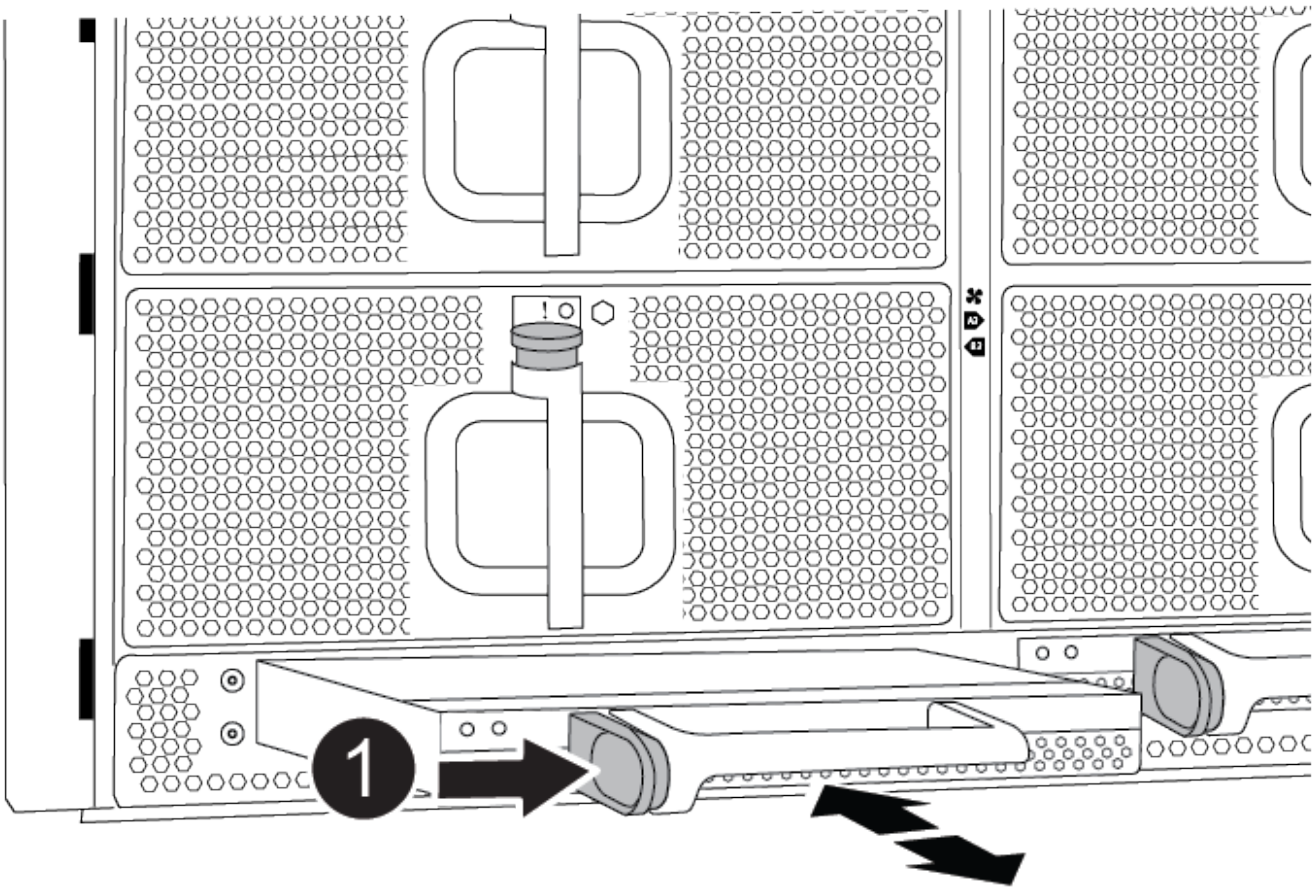

1	DCPM terra cotta locking button

3. Set the DCPM aside in a safe place and repeat this step for the remaining DCPM.

Step 6: Remove the USB LED module

Remove the USB LED modules.

Animation - Remove/install USB

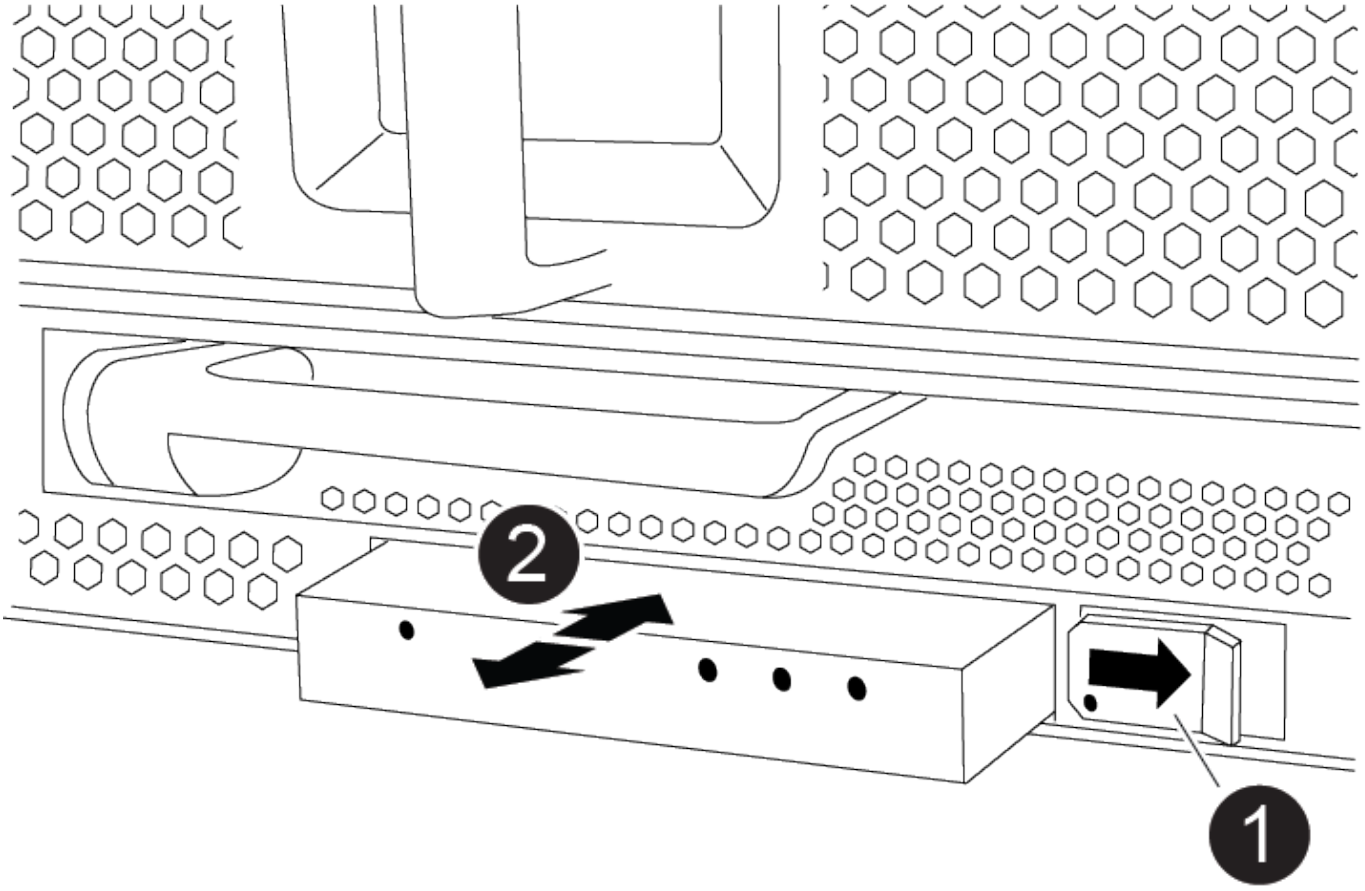

1	Eject the module.
2	Slide out of chassis.

1. Locate the USB LED module on the front of the impaired chassis, directly under the DCPM bays.
2. Press the black locking button on the right side of the module to release the module from the chassis, and then slide it out of the impaired chassis.
3. Set the module aside in a safe place.

Step 7: Remove chassis

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.

   If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

2. With the help of two or three people, slide the impaired chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.

4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.
7. Secure the rear of the chassis to the equipment rack or system cabinet.
8. If you are using the cable management brackets, remove them from the impaired chassis, and then install them on the replacement chassis.

Step 8: Install the de-stage controller power module

When the replacement chassis is installed into the rack or system cabinet, you must reinstall the de-stage controller power modules into it.

1. If you are not already grounded, properly ground yourself.
2. Align the end of the DCPM with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

3. Repeat this step for the remaining DCPM.

Step 9: Install fans into the chassis

To install the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

1. If you are not already grounded, properly ground yourself.
2. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

3. Repeat these steps for the remaining fan modules.
4. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Step 10: Install I/O modules

To install I/O modules, including the NVRAM modules from the impaired chassis, follow the specific sequence of steps.

You must have the chassis installed so that you can install the I/O modules into the corresponding slots in the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. After the replacement chassis is installed in the rack or cabinet, install the I/O modules into their corresponding slots in the replacement chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage, and then push the I/O cam latch all the way up to lock the module in place.

3. Recable the I/O module, as needed.
4. Repeat the preceding step for the remaining I/O modules that you set aside.

> If the impaired chassis has blank I/O panels, move them to the replacement chassis at this time.

Step 11: Install the power supplies

Installing the power supplies when replacing a chassis involves installing the power supplies into the replacement chassis, and connecting to the power source.

1. If you are not already grounded, properly ground yourself.
2. Make sure the power supplies rockers are in the off position.
3. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

> Do not use excessive force when sliding the power supply into the system. You can damage the connector.

4. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

> Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

5. Repeat the preceding steps for any remaining power supplies.

Step 12: Install the USB LED modules

Install the USB LED modules in the replacement chassis.

1. Locate the USB LED module slot on the front of the replacement chassis, directly under the DCPM bays.
2. Align the edges of the module with the USB LED bay, and gently push the module all the way into the chassis until it clicks into place.

Step 13: Install the controller

After you install the controller module and any other components into the replacement chassis, boot it.

1. If you are not already grounded, properly ground yourself.
2. Connect the power supplies to different power sources, and then turn them on.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

> Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the console to the controller module, and then reconnect the management port.
5. With the cam handle in the open position, slide the controller module into the chassis and firmly push the

controller module in until it meets the midplane and is fully seated, and then close the cam handle until it clicks into the locked position.

> Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.

The controller module begins to boot as soon as it is fully seated in the chassis.

6. Repeat the preceding steps to install the second controller into the replacement chassis.
7. Boot each controller.

Restore and verify the configuration - ASA A900

To complete the chassis replacement, you must complete specific tasks.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis *ha-state*`

      The value for HA-state can be one of the following:

      - ha
      - non-ha
3. Confirm that the setting has changed: `ha-config show`
4. If you have not already done so, recable the rest of your system.

Step 2: Bring up the system

1. If you have not done so, plug the power cables back into the PSUs.
2. Turn on the PSUs by toggling the rocker switched to **ON**, and wait for the controllers to power up completely.
3. Check the front and the back of the chassis and controllers for any fault lights after power up.
4. Connect to the SP or BMC IP address of the nodes via SSH. This will be the same address used to shut down the nodes.
5. Perform additional health checks as described in How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP
6. Turn AutoSupport back on (end the maintenance window message):
   `system node autosupport invoke -node * -type all -message MAINT=end`

As a best practice, you should do the following:

- Resolve any Active IQ Wellness Alerts and Risks (Active IQ will take time to process post-power up AutoSupports - expect a delay in results)
- Run Active IQ Config Advisor
- Check system health using How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Replace the controller module - ASA A900

To replace the impaired controller module, you must shut down the impaired controller, move the internal components to the replacement controller module, install the replacement controller module, and reboot the replacement controller.

Before you begin

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system has a V_StorageAttach license, you must refer to the additional required steps before performing this procedure.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the replacement controller so that the replacement controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The impaired controller is the controller that is being replaced.

    - The replacement controller is the new controller that is replacing the impaired controller.
    - The healthy controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - ASA A900

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the controller module hardware - ASA A900

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

The following animation shows the whole process of moving components from the impaired to the replacement controller.

Animation - Move components to replacement controller

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

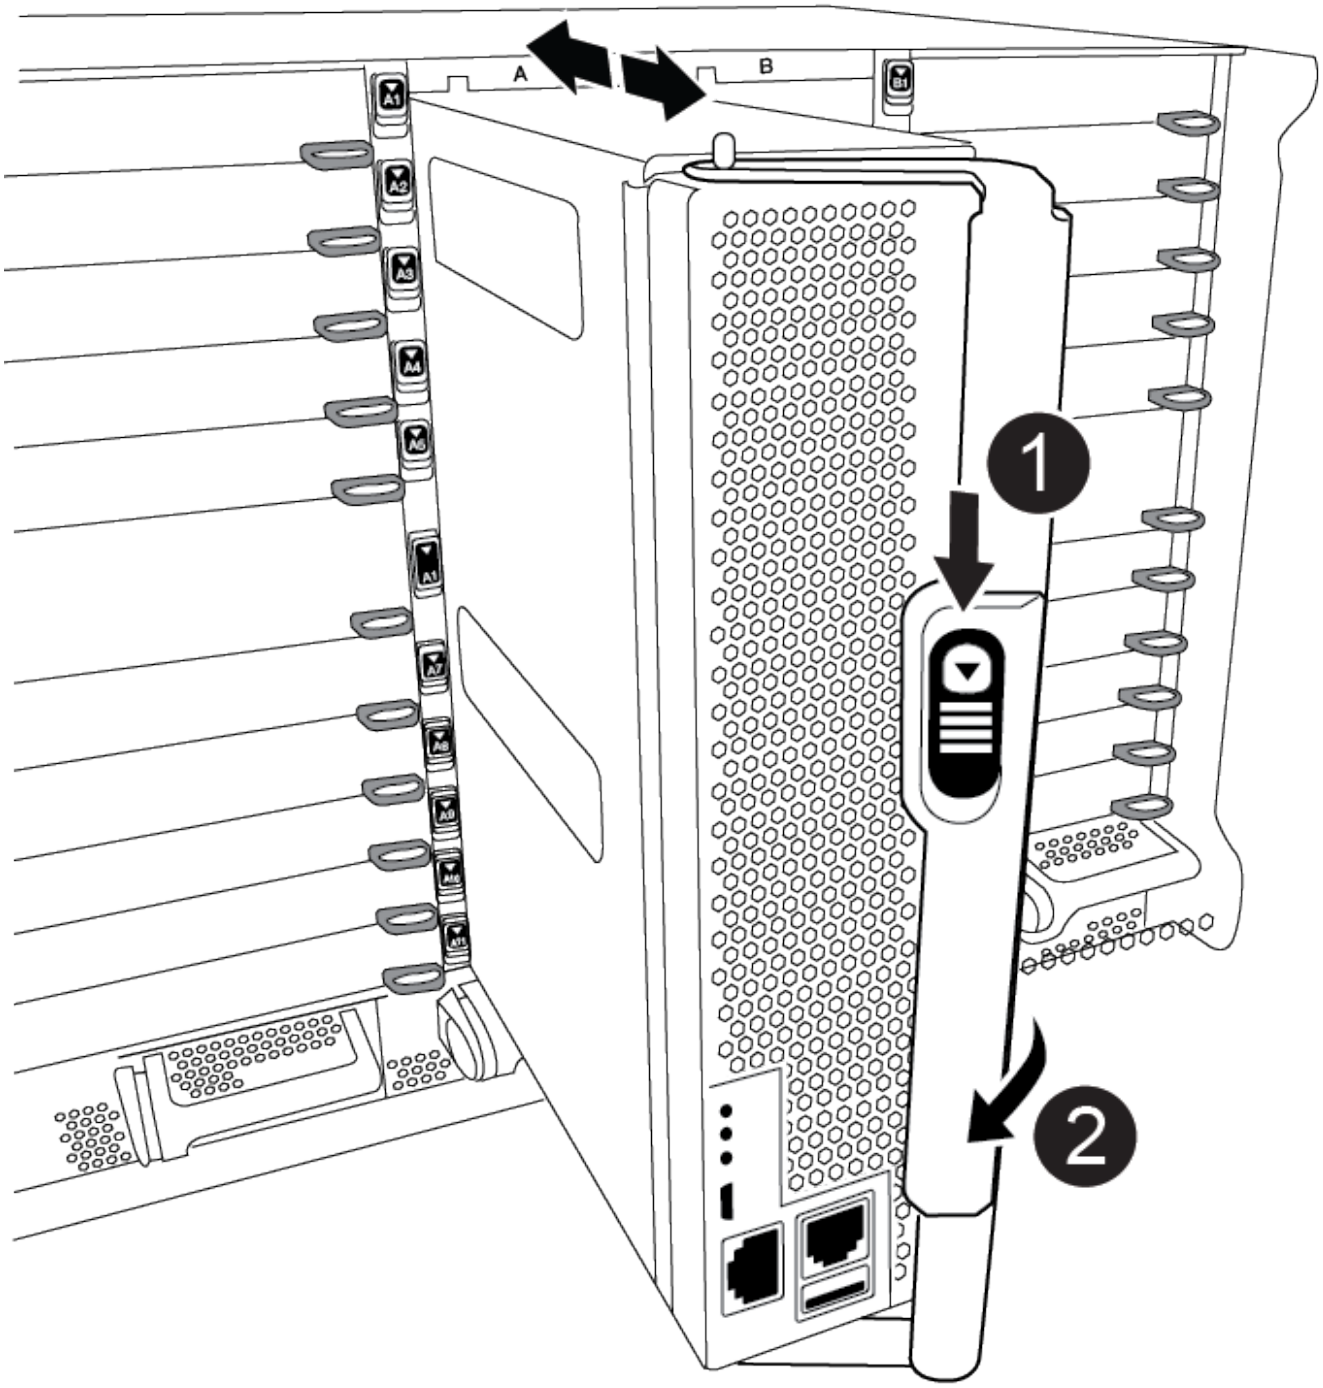

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

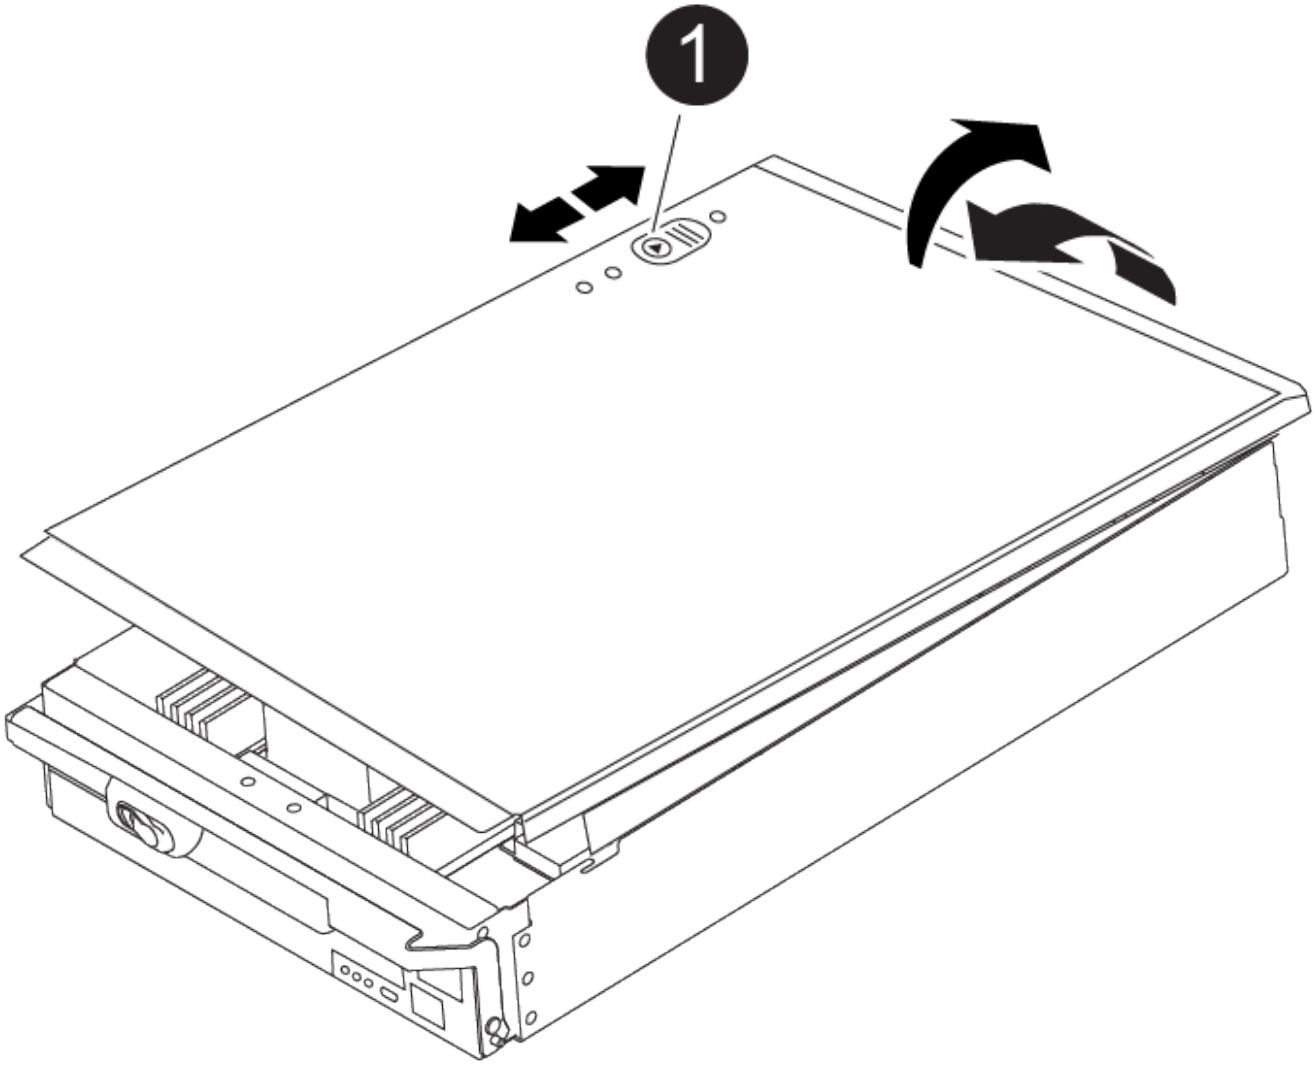

1	Controller module cover locking button

Step 2: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller and insert it in the new controller.

1. Locate the boot media using the following illustration or the FRU map on the controller module:

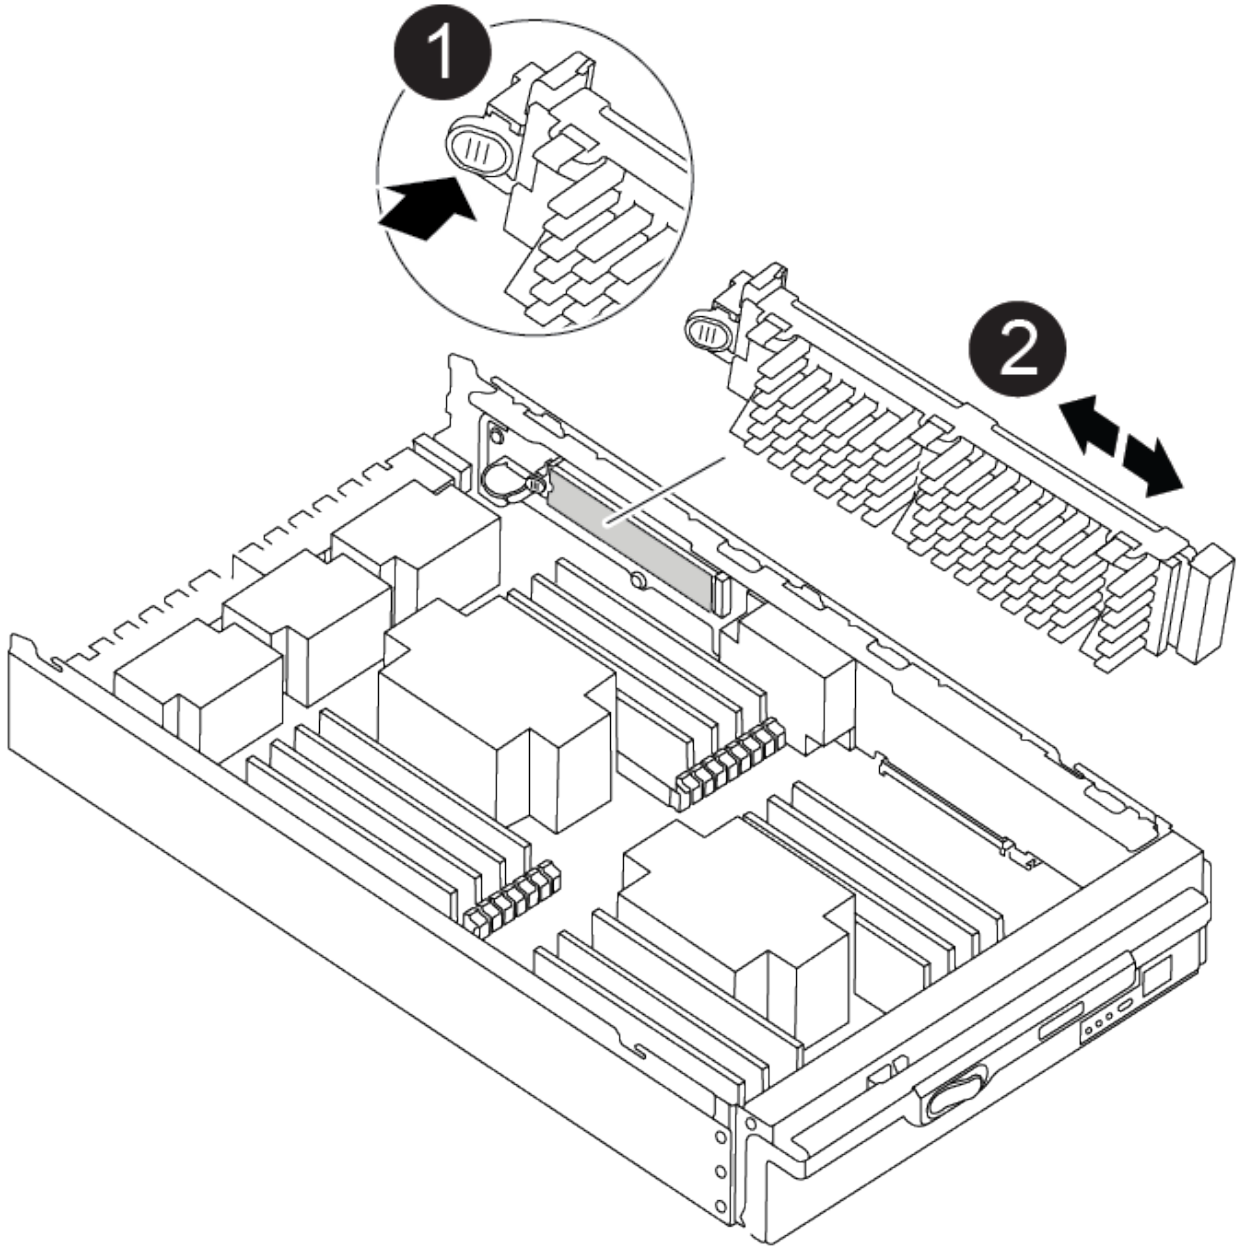

1	Press release tab
2	Boot media

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the system DIMMs

To move the DIMMs, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

> The VER2 controller has fewer DIMM sockets. There is no reduction in the number of DIMMs supported or change in the DIMM socket numbering. When moving the DIMMs to the new controller module, install the DIMMs into the same socket number/location as the impaired controller module. See the FRU map diagram on the VER2 controller module for DIMM socket locations.

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

> Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

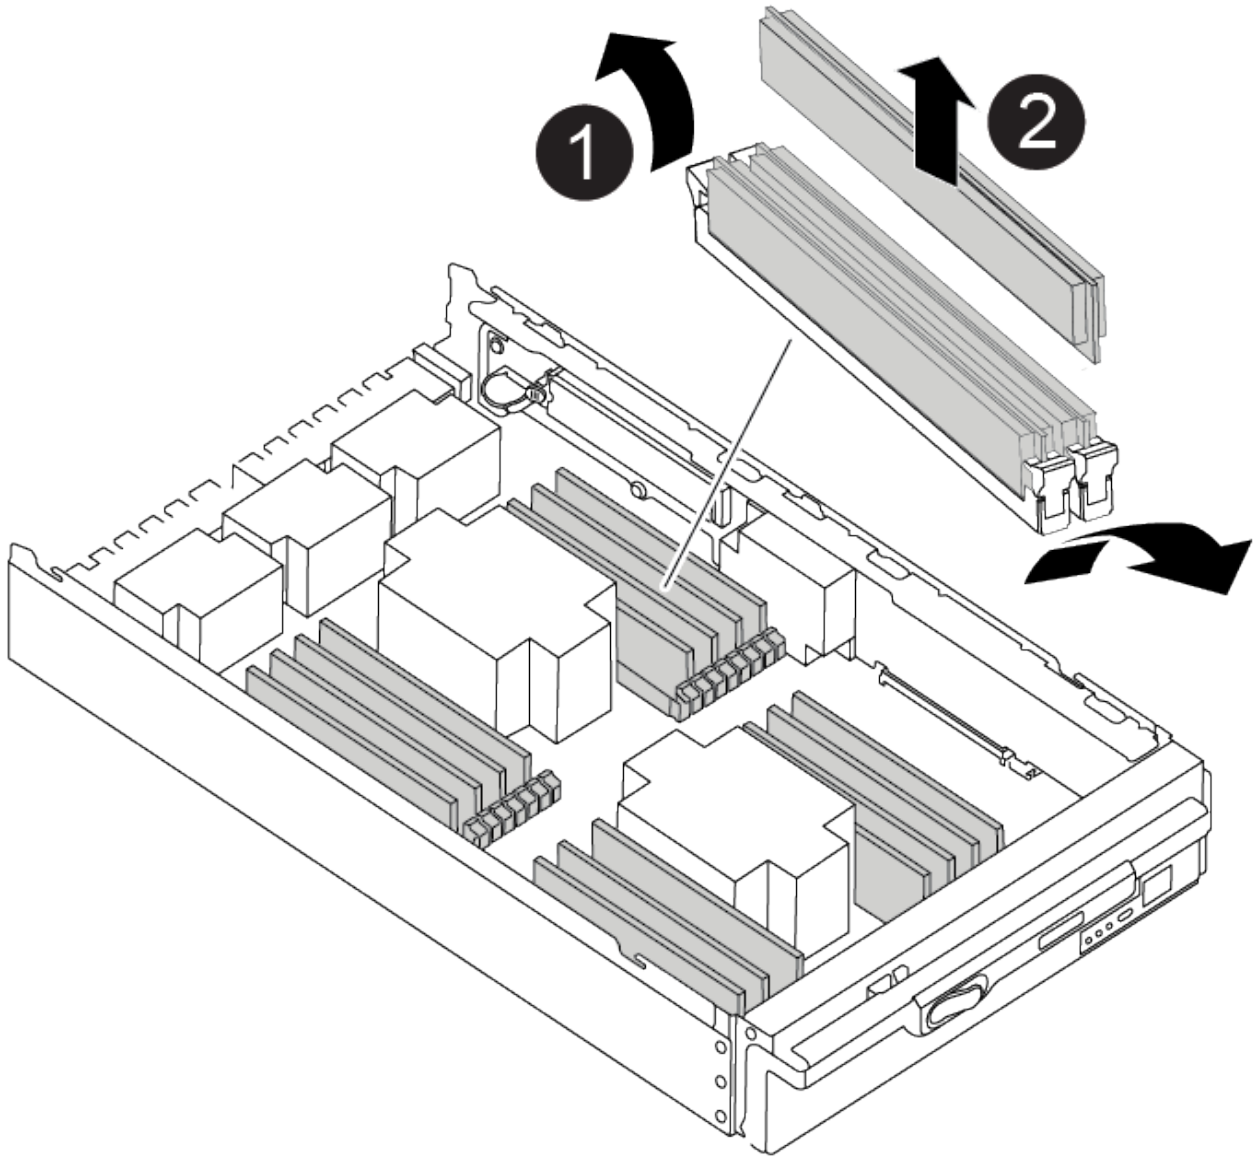

1	DIMM ejector tabs
2	DIMM

5. Locate the slot where you are installing the DIMM.
6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
9. Repeat these steps for the remaining DIMMs.

Step 4: Install the controller

After you install the components into the replacement controller module, you must install the replacement controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Animation - Install controller

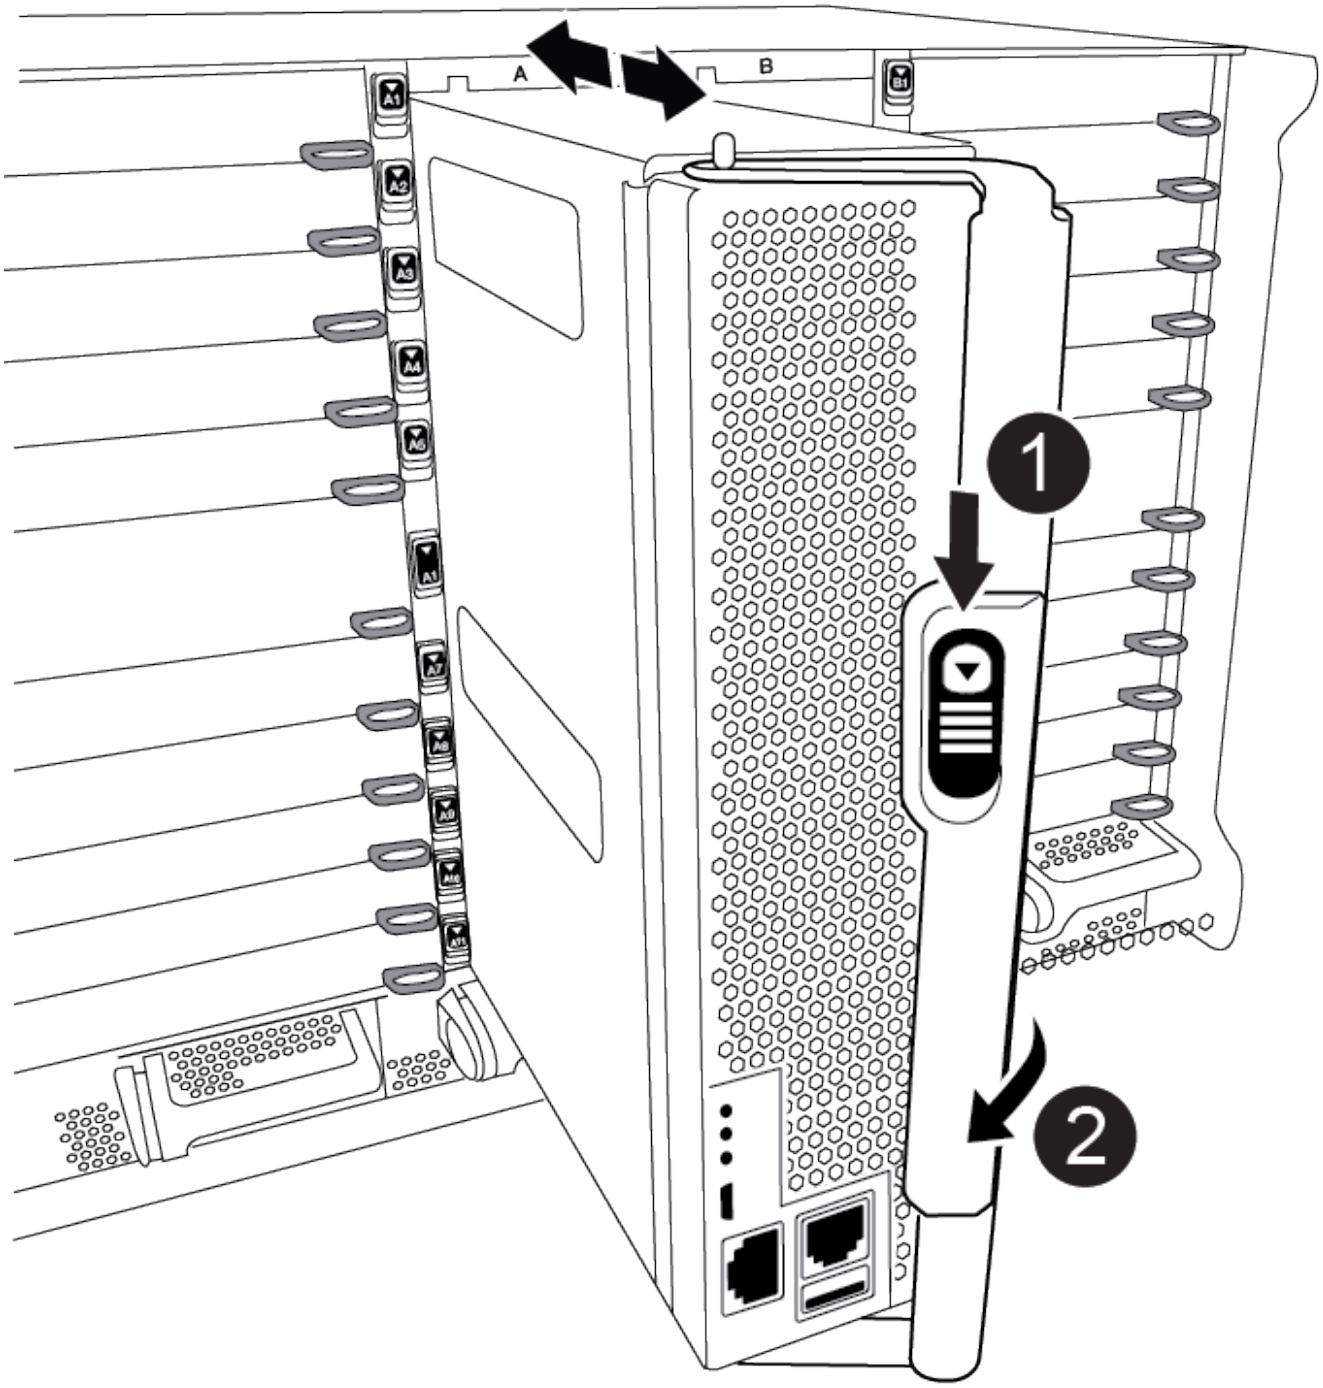

1	Cam handle release button
2	Cam handle

Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in

the following sections.

> You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   c. Rotate the controller module cam handle to the locked position.
   d. Interrupt the boot process by pressing `Ctrl-C` when you see Press Ctrl-C for Boot Menu.
   e. Select the option to boot to LOADER.

Restore and verify the system configuration - ASA A900

After completing the hardware replacement, you verify the low-level system configuration of the replacement controller, and reconfigure system settings as necessary.

Step 1: Set and verify the system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*

6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the controller's HA state

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the replacement controller module, verify that all components display the same `HA` state: `ha-config show`

If your system is in…	The HA state for all components should be…
An HA pair	ha
A MetroCluster FC configuration with four or more nodes	mcc
A MetroCluster IP configuration	mccip

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller` *`ha-state`*
3. If the displayed system state of the chassis does not match your system configuration, set the `HA` state for the chassis: `ha-config modify chassis` *`ha-state`*

Recable the system - ASA A900

Continue the replacement procedure by recabling the storage and network conigurations.

Step 1: Recable the system

You must recable the controller module's storage and network connections.

Steps

1. Recable the system.
2. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

> The system ID and disk assignment information reside in the NVRAM module, which is in a module separate from the controller module and not impacted by the controller module replacement.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`
   c. Wait for the savecore command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *`local-node-name`* `partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

- Restore onboard key management encryption keys
- Restore external key management encryption keys

6. Give back the controller:

   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The 'metrocluster node show -fields node-systemid' command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR

home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

  For more information, see Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration topic.

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - ASA A900

To complete the replacement procedure and restore your system to full operation, you must recable the storage, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Install licenses for the new controller

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

  Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The licenses keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.
- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

   

   The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - ASA A900

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

Before you begin

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

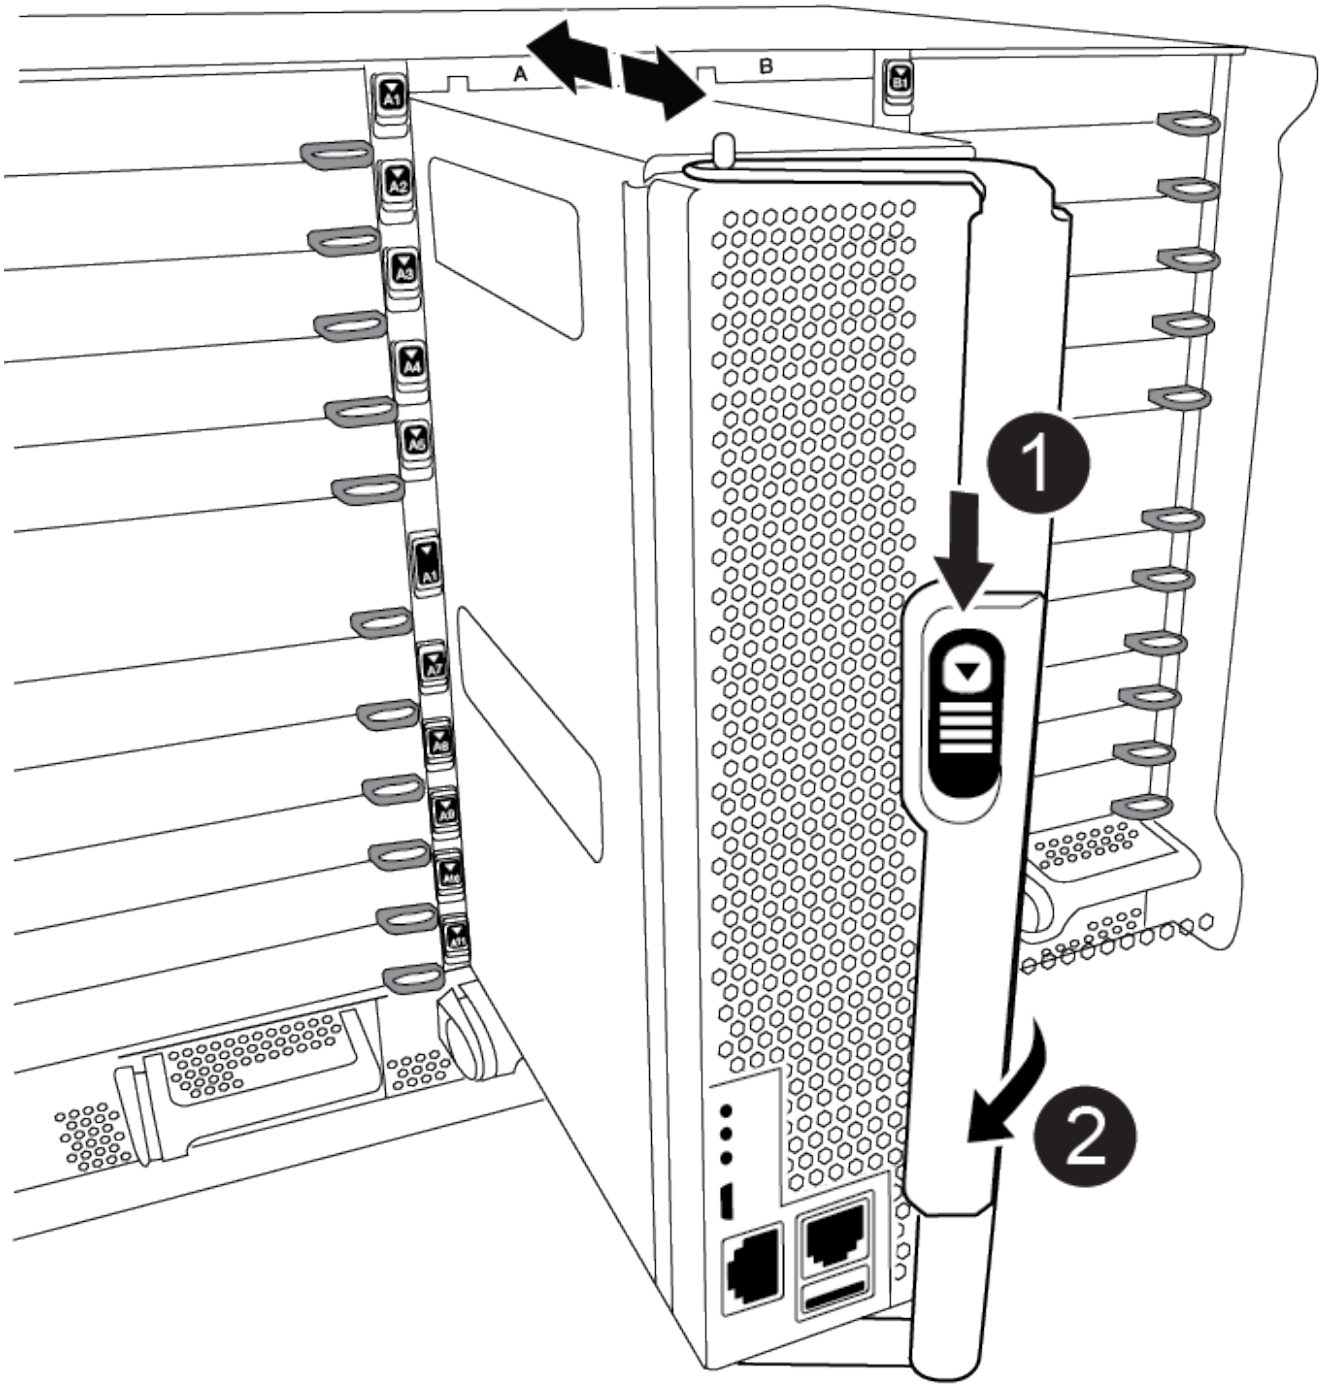

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

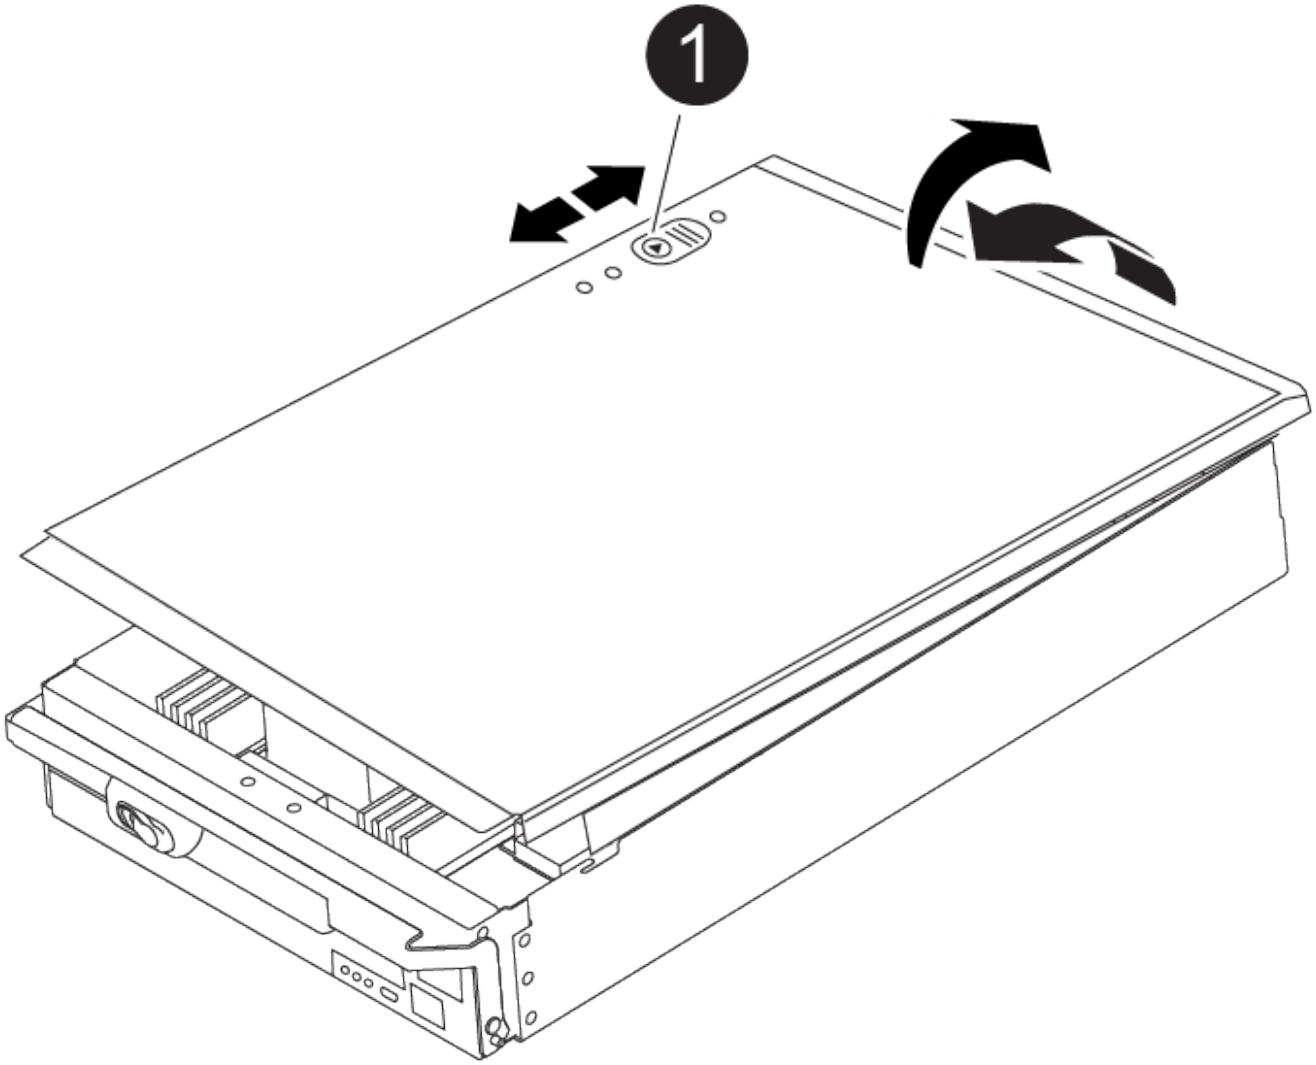

1	Controller module cover locking button

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

> The VER2 controller has fewer DIMM sockets. There is no reduction in the number of DIMMs supported or change in the DIMM socket numbering. When moving the DIMMs to the new controller module, install the DIMMs into the same socket number/location as the impaired controller module. See the FRU map diagram on the VER2 controller module for DIMM socket locations.

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.

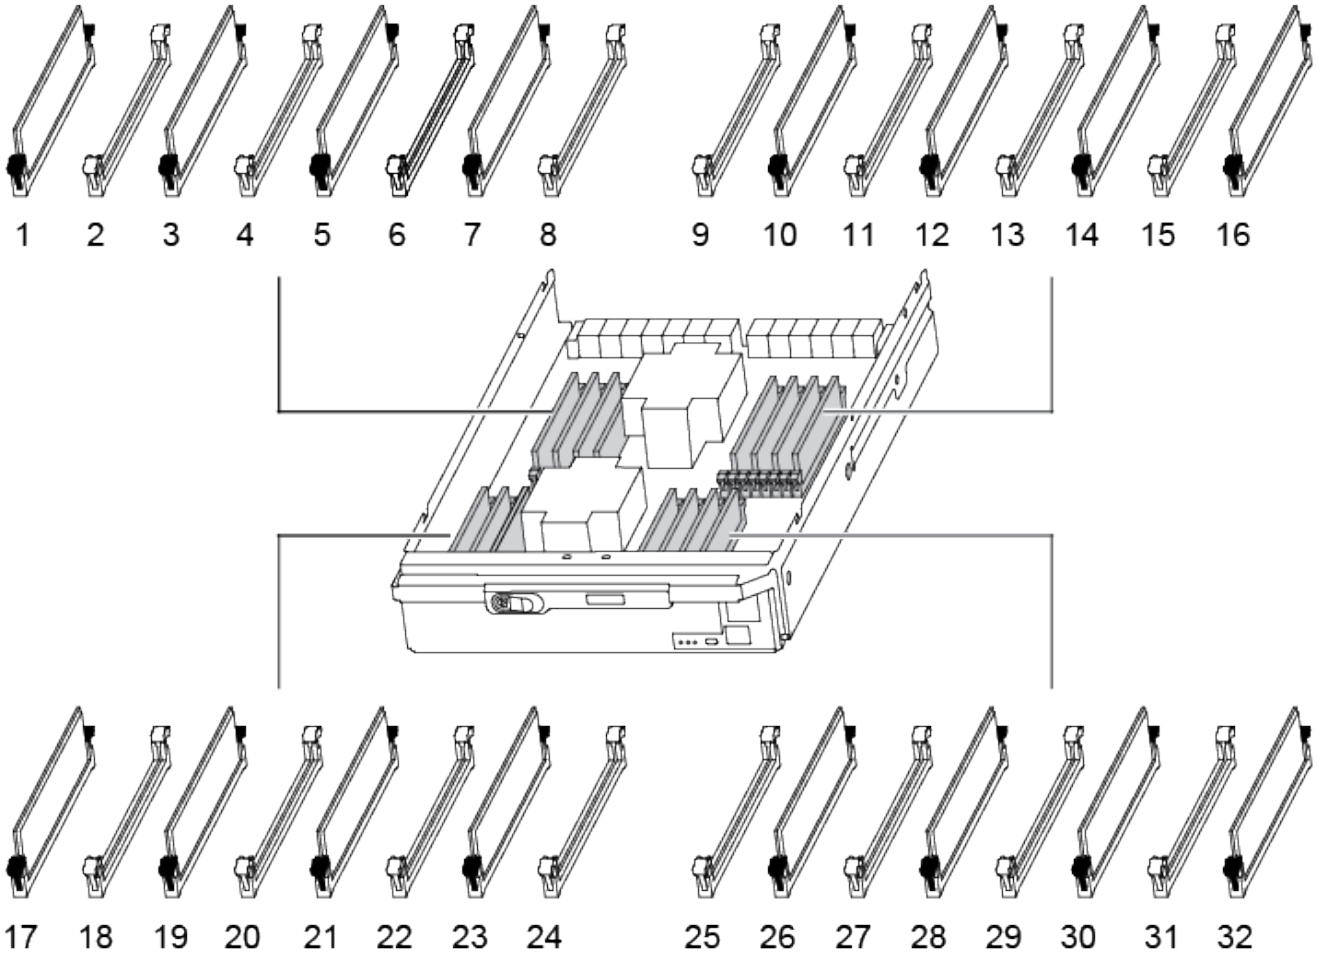

3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

Animation - Replace DIMM

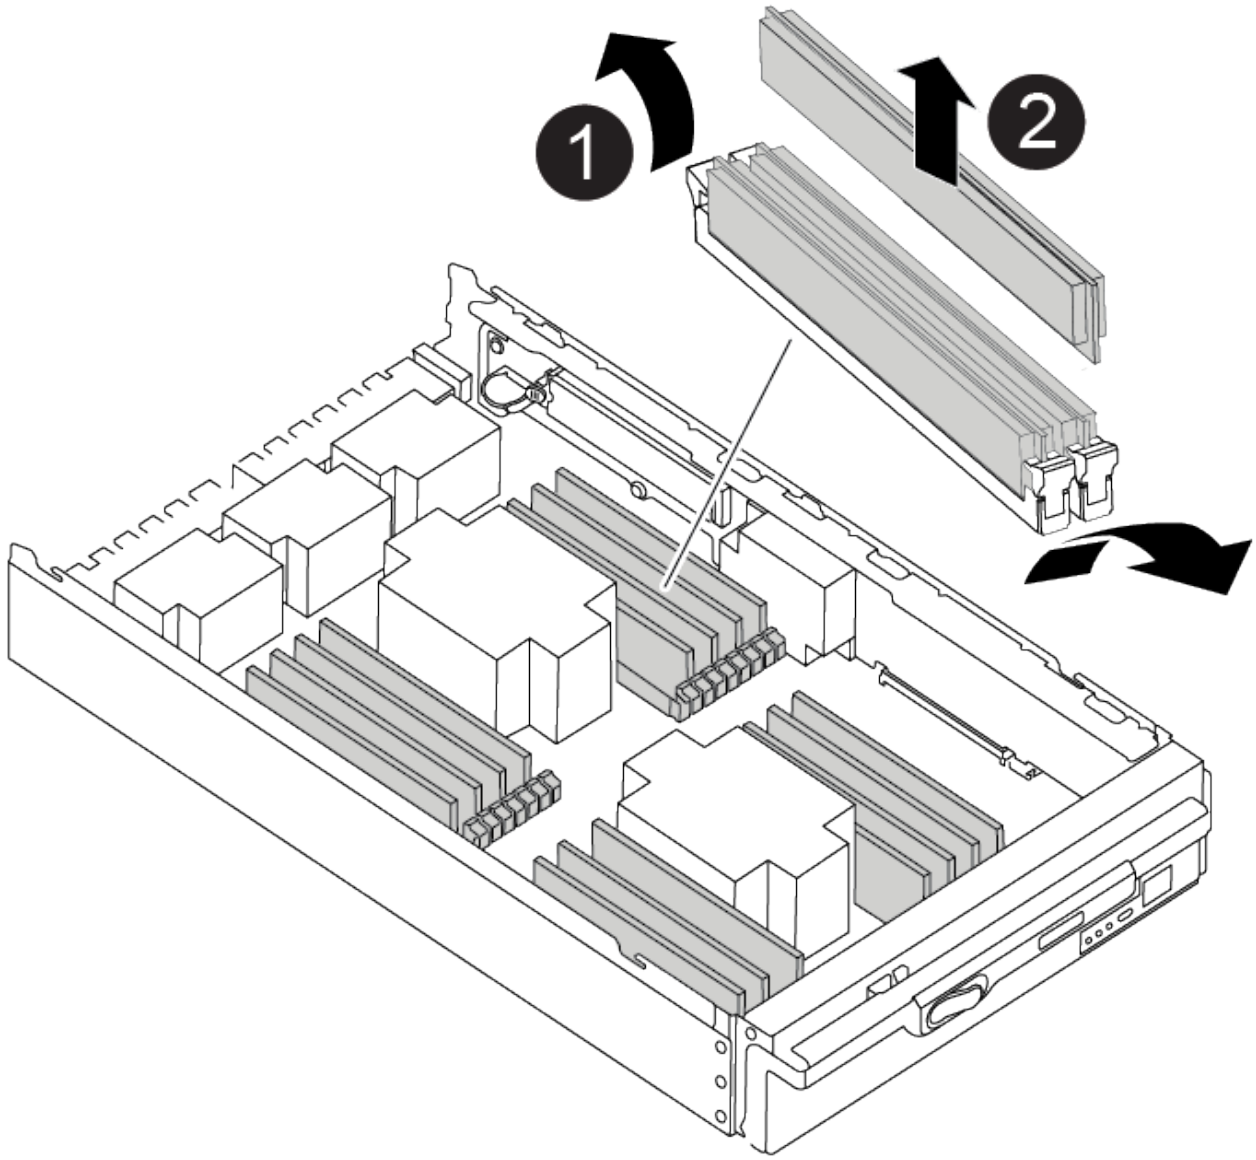

1	DIMM ejector tabs
2	DIMM

4. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

5. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Close the controller module cover.

Step 4: Install the controller

After you install the components into the controller module, you must install the controller module back into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.

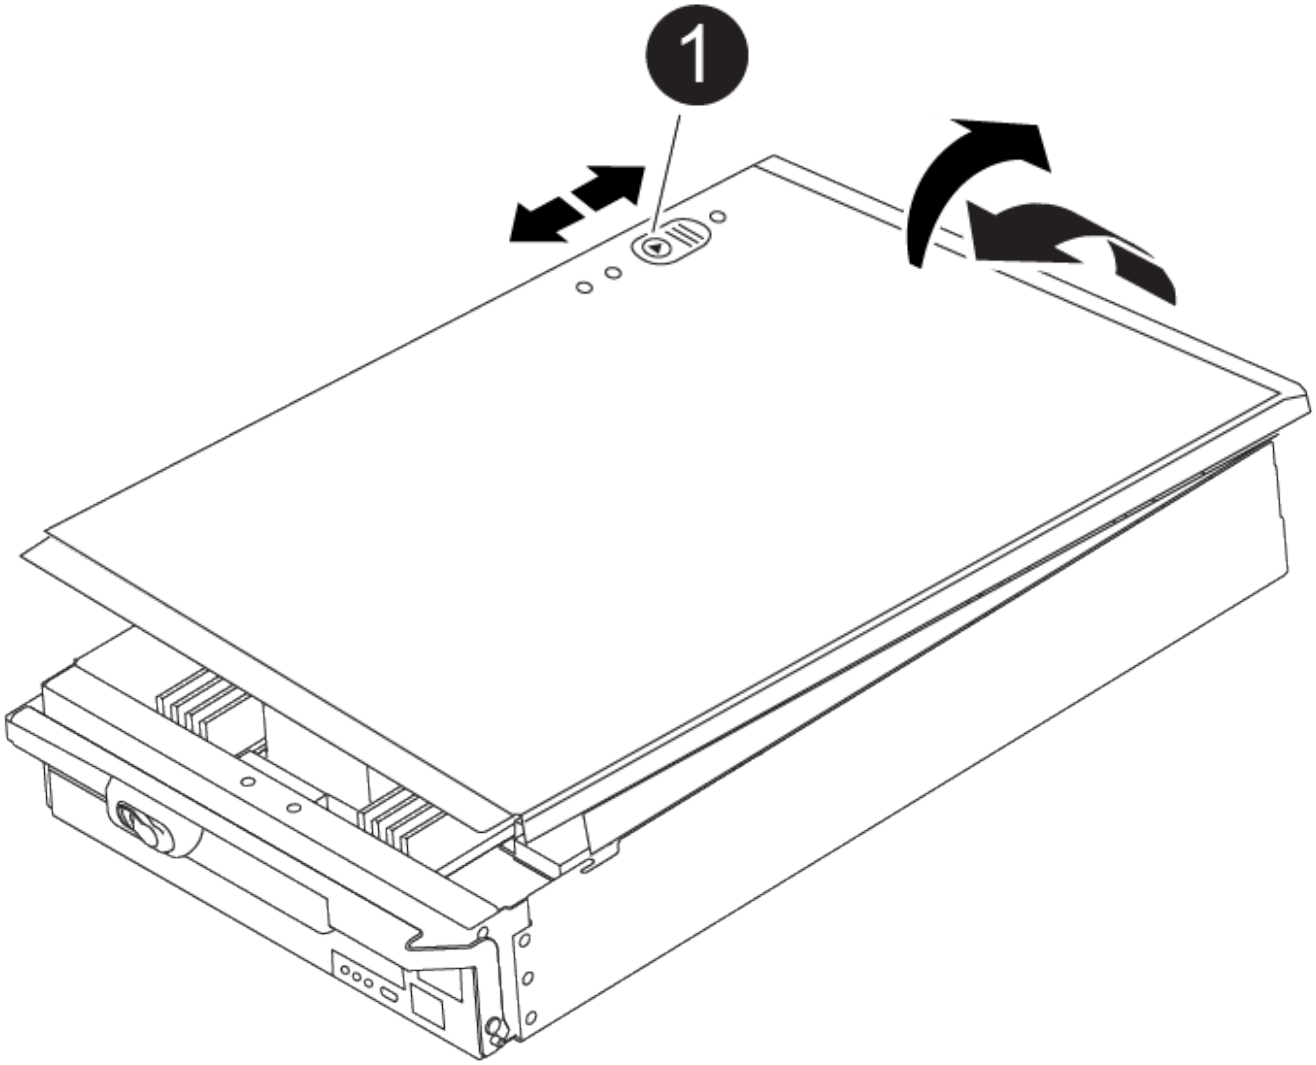

1	Controller module cover locking button

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Animation - Install controller

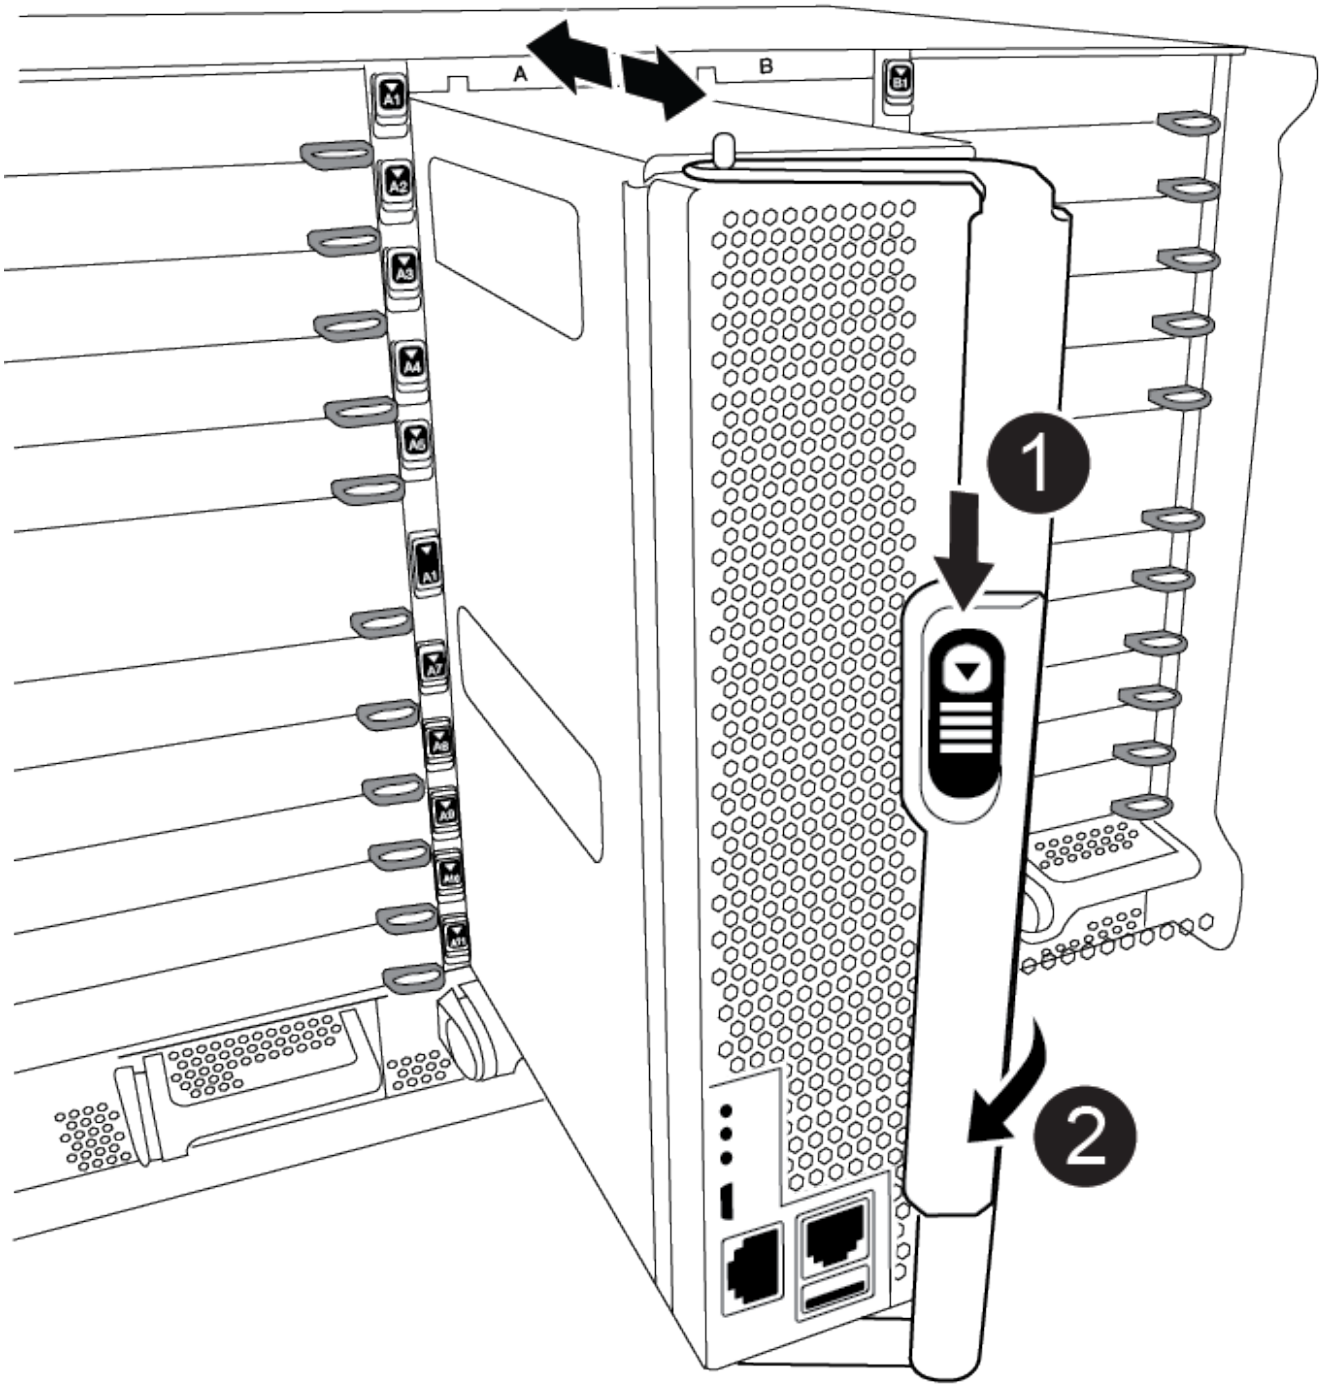

1	Cam handle release button
2	Cam handle

Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   d. Interrupt the boot process by pressing `Ctrl-C` when you see Press Ctrl-C for Boot Menu.
   e. Select the option to boot to Maintenance mode from the displayed menu.

Step 5: Run system-level diagnostics

After installing a new DIMM, you should run diagnostics.

Your system must be at the LOADER prompt to start System Level Diagnostics.

All commands in the diagnostic procedures are issued from the controller where the component is being replaced.

1. If the controller to be serviced is not at the LOADER prompt, perform the following steps:
   a. Select the Maintenance mode option from the displayed menu.
   b. After the controller boots to Maintenance mode, halt the controller: `halt`

      After you issue the command, you should wait until the system stops at the LOADER prompt.

      During the boot process, you can safely respond `y` to prompts.

      - If a prompt appears warning that when entering Maintenance mode in an HA configuration, you must ensure that the healthy controller remains down.

2. At the LOADER prompt, access the special drivers specifically designed for system-level diagnostics to function properly: `boot_diags`

   During the boot process, you can safely respond `y` to the prompts until the Maintenance mode prompt (*>) appears.

3. Run diagnostics on the system memory: `sldiag device run -dev mem`
4. Verify that no hardware problems resulted from the replacement of the DIMMs: `sldiag device status`

`-dev mem -long -state failed`

System-level diagnostics returns you to the prompt if there are no test failures, or lists the full status of failures resulting from testing the component.

5. Proceed based on the result of the preceding step:

If the system-level diagnostics tests…	Then…
Were completed without any failures	a. Clear the status logs: `sldiag device clearstatus` b. Verify that the log was cleared: `sldiag device status` The following default response is displayed: SLDIAG: No log messages are present. c. Exit Maintenance mode: `halt` The controller displays the LOADER prompt. d. Boot the controller from the LOADER prompt: `bye` e. Return the controller to normal operation:

If your controller is in…	Then…
An HA pair	Perform a give back: `storage failover giveback -ofnode replacement_node_name` **Note:** If you disabled automatic giveback, re-enable it with the storage failover modify command.

If your controller is in…	Then…
Resulted in some test failures	Determine the cause of the problem: a. Exit Maintenance mode: `halt` After you issue the command, wait until the system stops at the LOADER prompt. b. Verify that you have observed all the considerations identified for running system-level diagnostics, that cables are securely connected, and that hardware components are properly installed in the storage system. c. Boot the controller module you are servicing, interrupting the boot by pressing `Ctrl-C` when prompted to get to the Boot menu: ◦ If you have two controller modules in the chassis, fully seat the controller module you are servicing in the chassis. The controller module boots up when fully seated. ◦ If you have one controller module in the chassis, connect the power supplies, and then turn them on. d. Select Boot to maintenance mode from the menu. e. Exit Maintenance mode by entering the following command: `halt` After you issue the command, wait until the system stops at the LOADER prompt. f. Rerun the system-level diagnostic test.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the DCPM containing the NVRAM11 battery - ASA A900

To hot-swap a destage controller power module (DCPM), which contains the NVRAM11 battery, you must locate the failed DCPM module, remove it from the chassis, and install the replacement DCPM module.

You must have a replacement DCPM module in-hand before removing the failed module from the chassis and it must be replaced within five minutes of removal. Once the DCPM module is removed from the chassis, there is no shutdown protection for the controller module that owns the DCPM module, other than failover to the other controller module.

Step 1: Replace the DCPM module

To replace the DCPM module in your system, you must remove the failed DCPM module from the system and then replace it with a new DCPM module.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel on the front of the system and set it aside.
3. Locate the failed DCPM module in the front of the system by looking for the Attention LED on the module.

   The LED will be steady amber if the module is faulty.

   > The DCPM module must be replaced in the chassis within five minutes of removal or the associated controller will shut down.

4. Press the terra cotta release button on the module handle, and then slide the DCPM module out of the chassis.

   Animation - Remove/install DCPM

   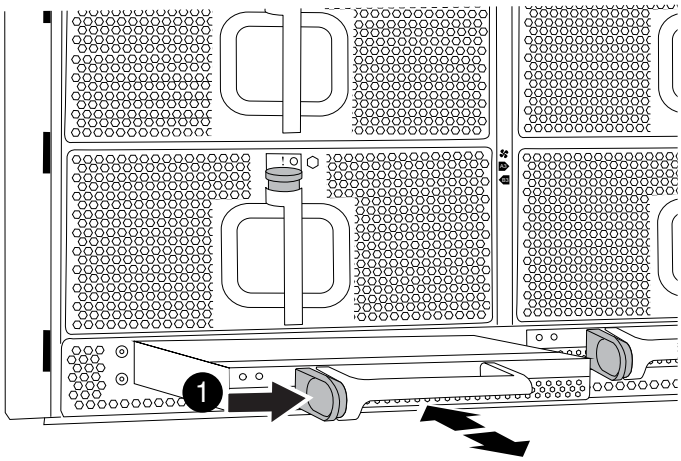

1	DCPM module terra cotta locking button

5. Align the end of the DCPM module with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   > The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

   The Amber LED flashes four times upon insertion and the green LED also flashes if the battery is providing a voltage. If it does not flash, it will likely need to be replaced.

Step 2: Dispose of batteries

You must dispose of batteries according to the local regulations regarding battery recycling or disposal. If you cannot properly dispose of batteries, you must return the batteries to NetApp, as described in the RMA instructions that are shipped with the kit.

Safety Information and Regulatory Notices

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a fan - ASA A900

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press the terra cotta button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

> The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

Animation - Remove/install fan

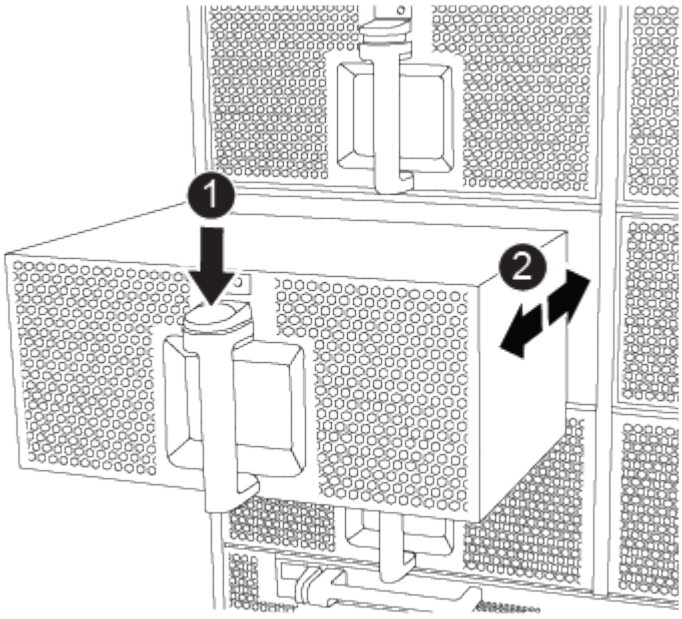

1	Terra cotta release button

2	Slide fan in/out of chassis

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Replace the I/O module - ASA A900

To replace an I/O module, you must perform a specific sequence of tasks.

- You can use this procedure with all versions of ONTAP supported by your system.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired node

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster

Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace I/O modules

To replace an I/O module, locate it within the chassis and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

3. Remove the target I/O module from the chassis:

   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

Make sure that you keep track of which slot the I/O module was in.

Animation - Remove/install I/O module

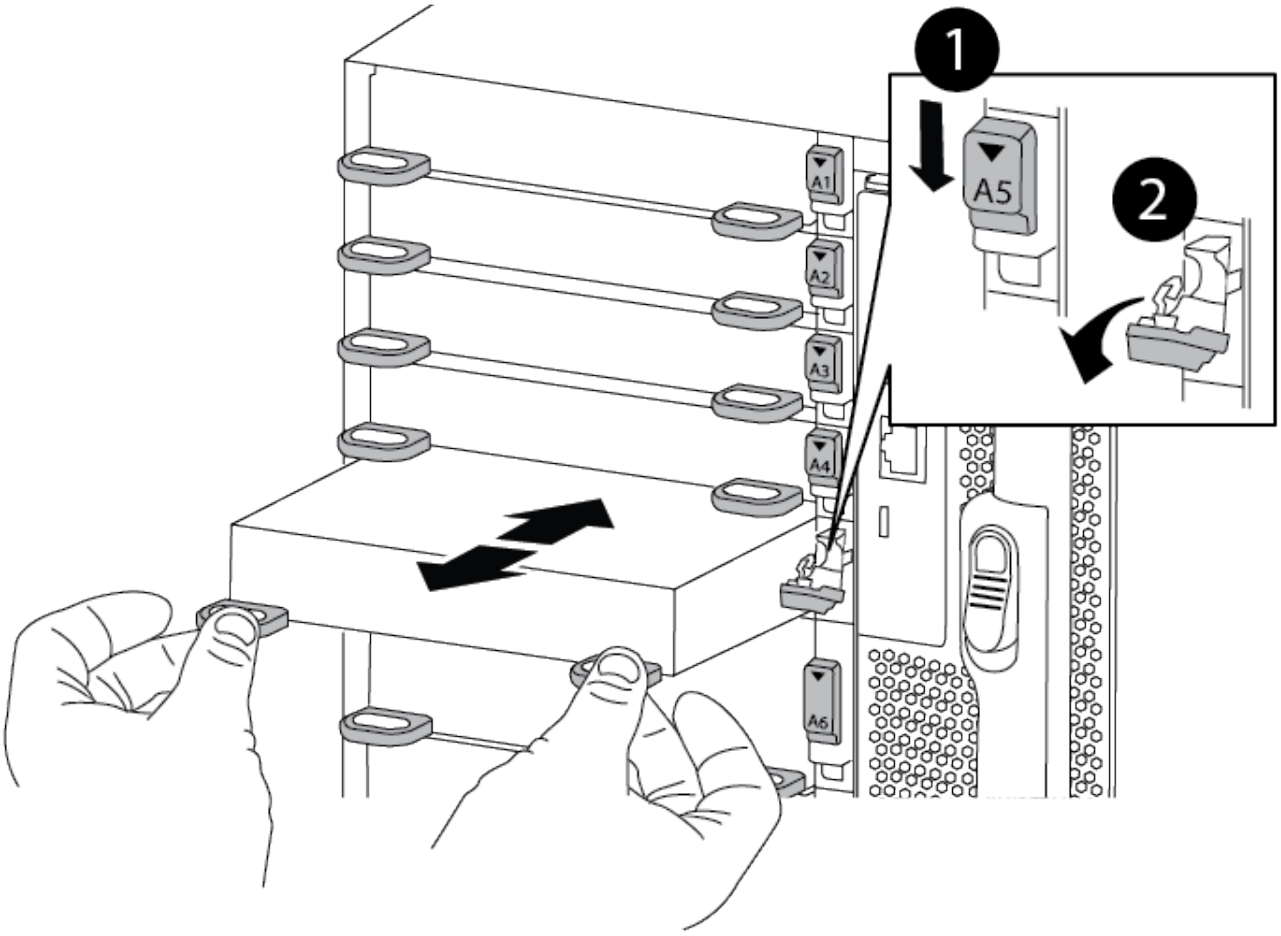

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Install the replacement I/O module into the chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.
6. Recable the I/O module, as needed.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller module.

If the new I/O module is not the same model as the failed module, you must first reboot the BMC.

Steps

1. Reboot the BMC if the replacement module is not the same model as the old module:
   a. From the LOADER prompt, change to advanced privilege mode: `priv set advanced`
   b. Reboot the BMC: `sp reboot`
2. From the LOADER prompt, reboot the node: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.
3. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs, convert these ports to 10 GbE connections by using the nicadmin convert command from Maintenance mode. See Convert 40GbE NIC ports into multiple 10GbE ports for 10GbE connectivity for more information.

   Be sure to exit Maintenance mode after completing the conversion.
4. Return the node to normal operation: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Add an I/O module - ASA A900

If the storage system has empty slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- To non-disruptively add an I/O module, you must take over the target controller, remove the slot blanking cover in the target slot or remove an existing I/O module, add the new or replacement I/O module, and then giveback the target controller.
- Make sure that all other components are functioning properly.

Add I/O module to an empty slot

You can add a new I/O module into a storage system with available empty slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the target slot blanking cover:
   a. Depress the lettered and numbered cam latch.
   b. Rotate the cam latch down until it is the open position.
   c. Remove the blanking cover.
3. Install the I/O module:
   a. Align the I/O module with the edges of the slot.
   b. Slide the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
4. If the replacement I/O module is a NIC, cable the module to the data switches.

   Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

5. Reboot the controller from the LOADER prompt: *bye*

   This reinitializes the PCIe cards and other components and reboots the node.

6. Give back the controller from the partner controller. `storage failover giveback -ofnode target_node_name`
7. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto -giveback true`
8. If you are using slots 3 and/or 7 for networking, use the `storage port modify -node <node name> -port <port name> -mode network` command to convert the slot for networking use.
9. Repeat these steps for controller B.
10. If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam latch.

      The cam latch moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

      Animation - Remove or replacing an I/O module

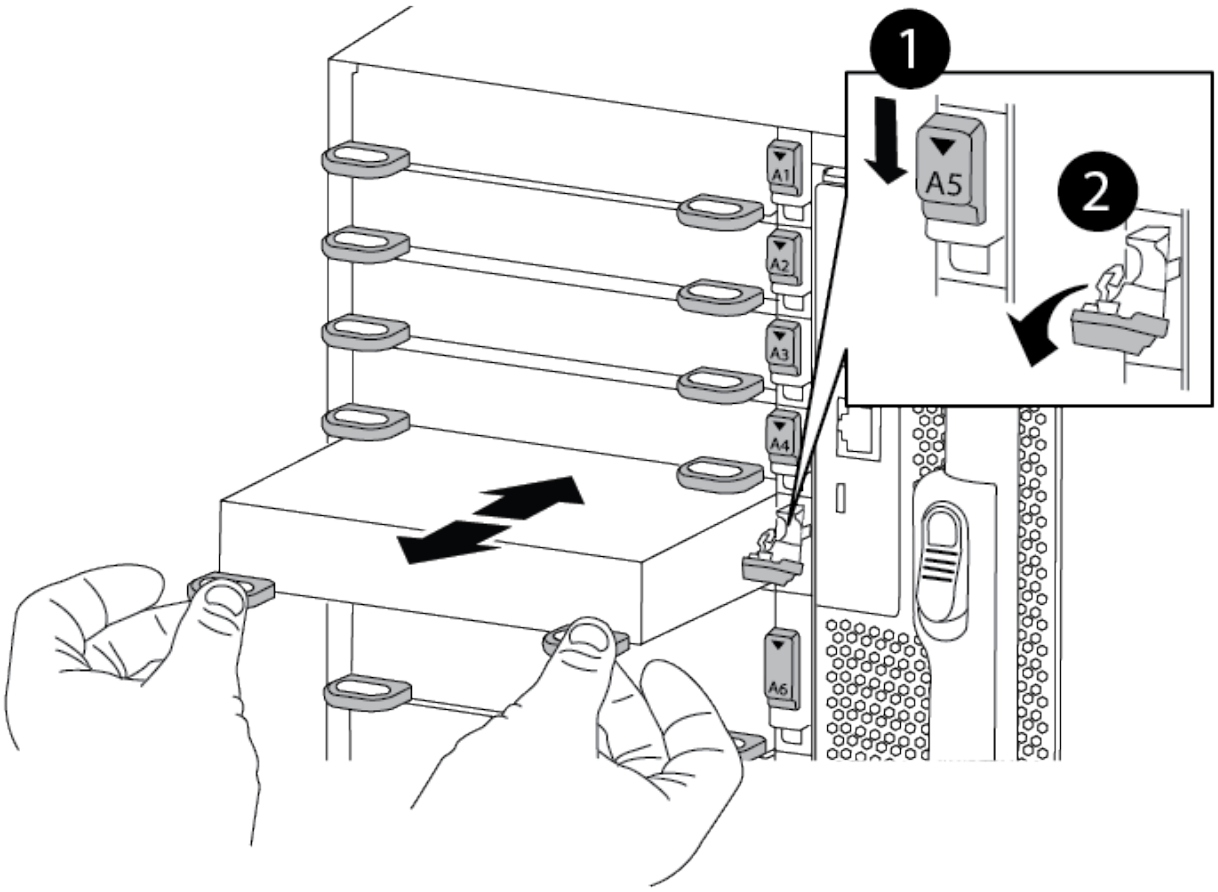

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Install the I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Slide the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
5. Repeat the remove and install steps to replace additional modules for controller A.
6. If the replacement I/O module is a NIC, cable the module or modules to the data switches.
7. Reboot the controller from the LOADER prompt:
   a. Check the version of BMC on the controller: `system service-processor show`
   b. Update the BMC firmware if needed: `system service-processor image update`
   c. Reboot the node: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.

If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement

8. Give back the controller from the partner controller. `storage failover giveback -ofnode target_node_name`
9. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto-giveback true`
10. If you added:

If I/O module is a…	Then…
NIC module in slots 3 or 7,	Use the `storage port modify -node *<node name> -port *<port name> -mode network` command for each port.
Storage module	Install and cable your NS224 shelves, as described in Hot-add workflow.

11. Repeat these steps for controller B.

Replace an LED USB module - ASA A900

The LED USB module provides connectivity to console ports and system status. Replacement of this module does not require tools and does not interrupt service.

Step 1: Replace the LED USB module

Steps

1. Remove the impaired LED USB module:

   Animation - Remove/install LED-USB module

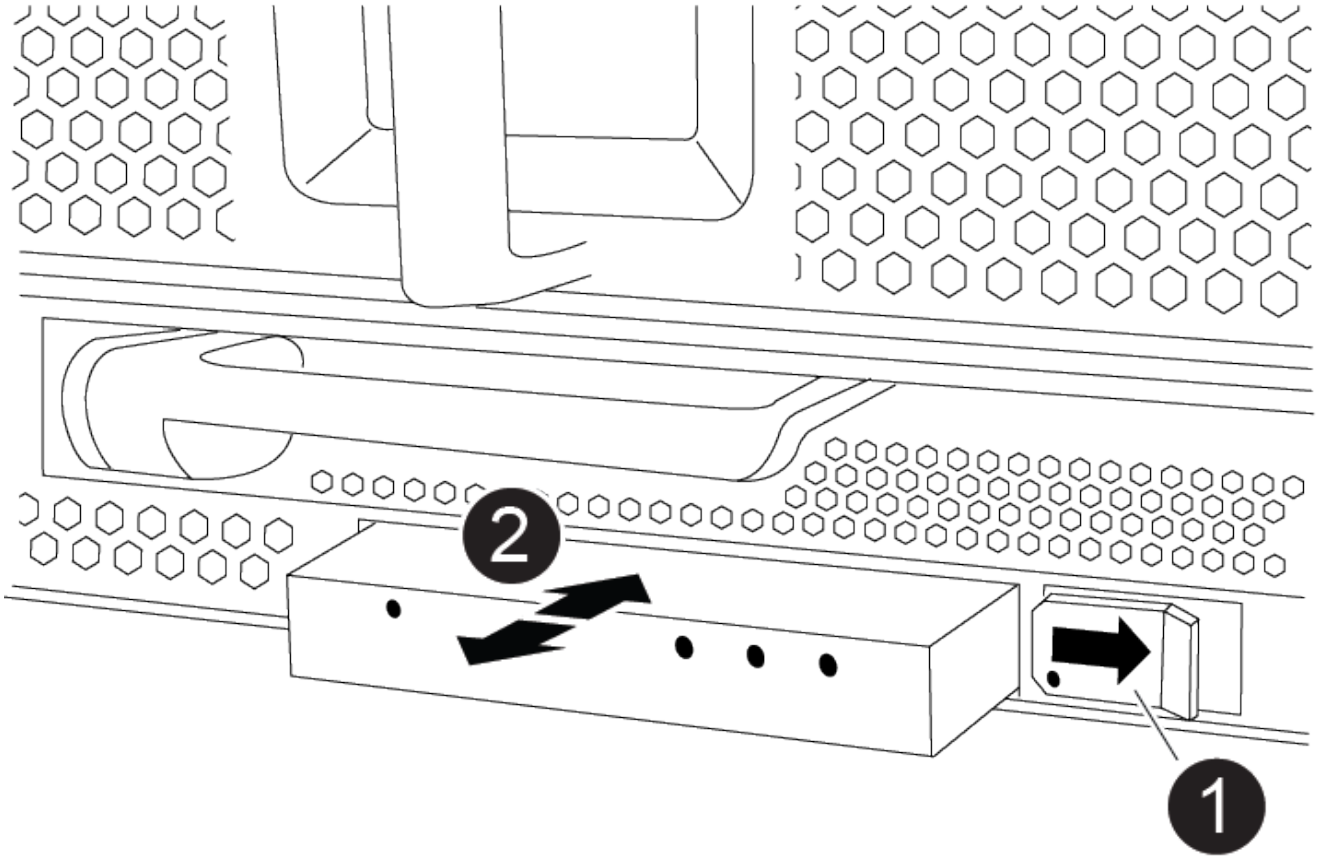

1	Locking button
2	USB LED module

a. With the bezel removed, locate the LED USB module at the front of the chassis, on the bottom left side.

b. Slide the latch to partially eject the module.

c. Pull the module out of the bay to disconnect it from the midplane. Do not leave the slot empty.

2. Install the new LED USB module:

a. Align the module to the bay with the notch in the corner of the module positioned near the slider latch on the chassis. The bay will prevent you from installing the module upside down.

b. Push the module into the bay until it is fully seated flush with the chassis.

There is an audible click when the module is secure and connected to the midplane.

Step 2: Return the failed component

1. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVRAM module and NVRAM DIMMs - ASA A900

The NVRAM module consists of the NVRAM11 and DIMMs. You can replace a failed

NVRAM module or the DIMMs inside the NVRAM module. To replace a failed NVRAM module, you must remove it from the chassis, move the DIMMs to the replacement module, and install the replacement NVRAM module into the chassis.

To replace and NVRAM DIMM, you must remove the NVRAM module from the chassis, replace the failed DIMM in the module, and then reinstall the NVRAM module.

About this task

Because the system ID is derived from the NVRAM module, if replacing the module, disks belonging to the system are reassigned to a new system ID.

Before you begin

- All disk shelves must be working properly.
- If your system is in an HA pair, the partner controller must be able to take over the controller associated with the NVRAM module that is being replaced.
- This procedure uses the following terminology:
  - The impaired controller is the controller on which you are performing maintenance.
  - The healthy controller is the HA partner of the impaired controller.
- This procedure includes steps for automatically reassigning disks to the controller module associated with the new NVRAM module. You must reassign the disks when directed to in the procedure. Completing the disk reassignment before giveback can cause issues.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You cannot change any disks or disk shelves as part of this procedure.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 6 in the chassis and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the target NVRAM module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The NVRAM module disengages from the chassis and moves out a few inches.

   c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

      Animation - Replace the NVRAM module

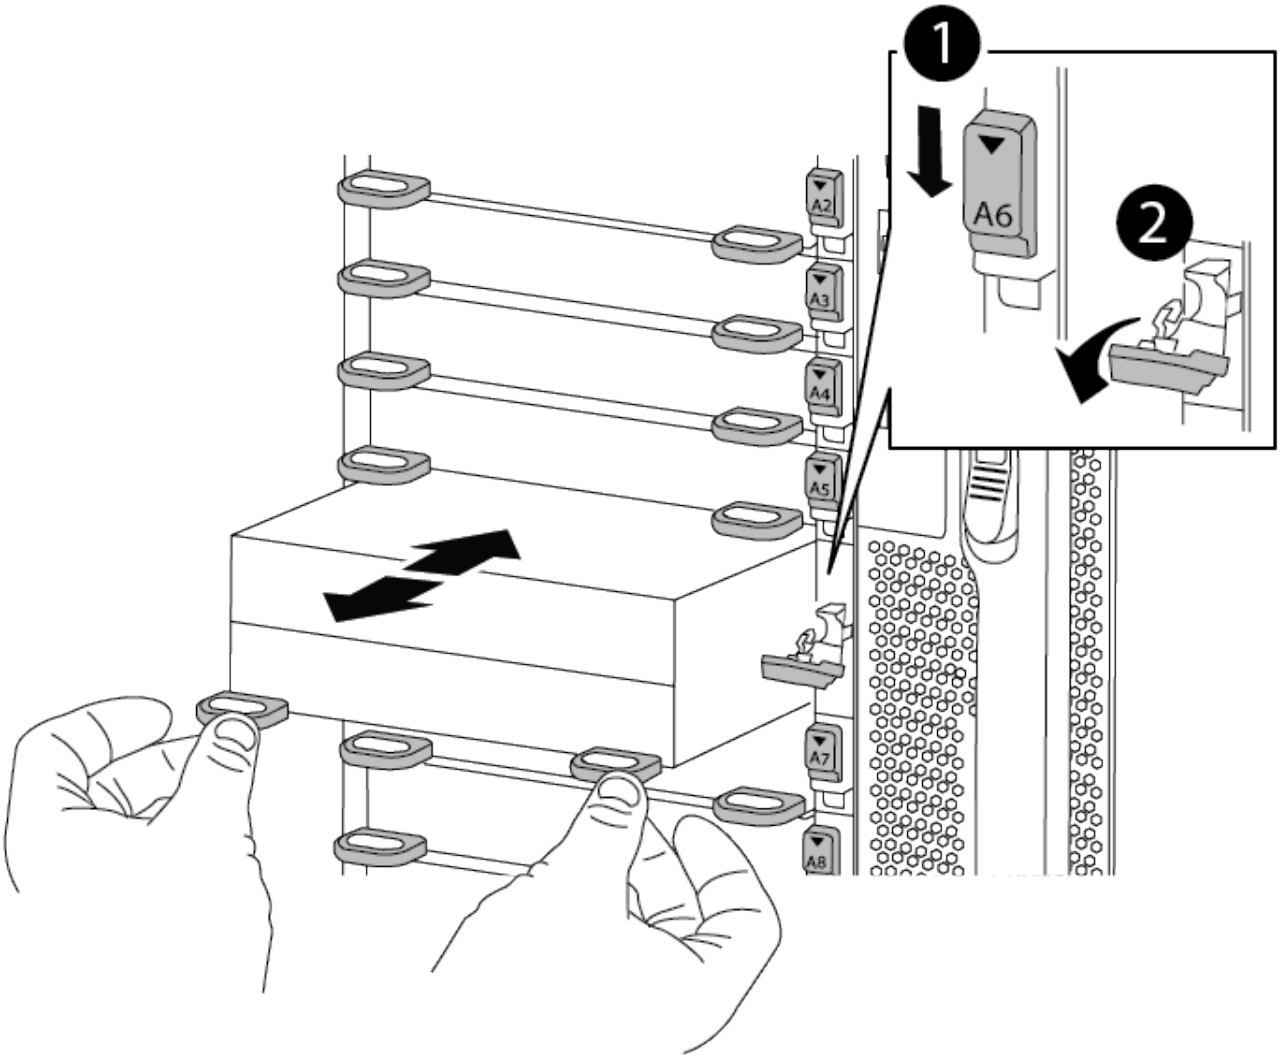

1	Lettered and numbered cam latch
2	Cam latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

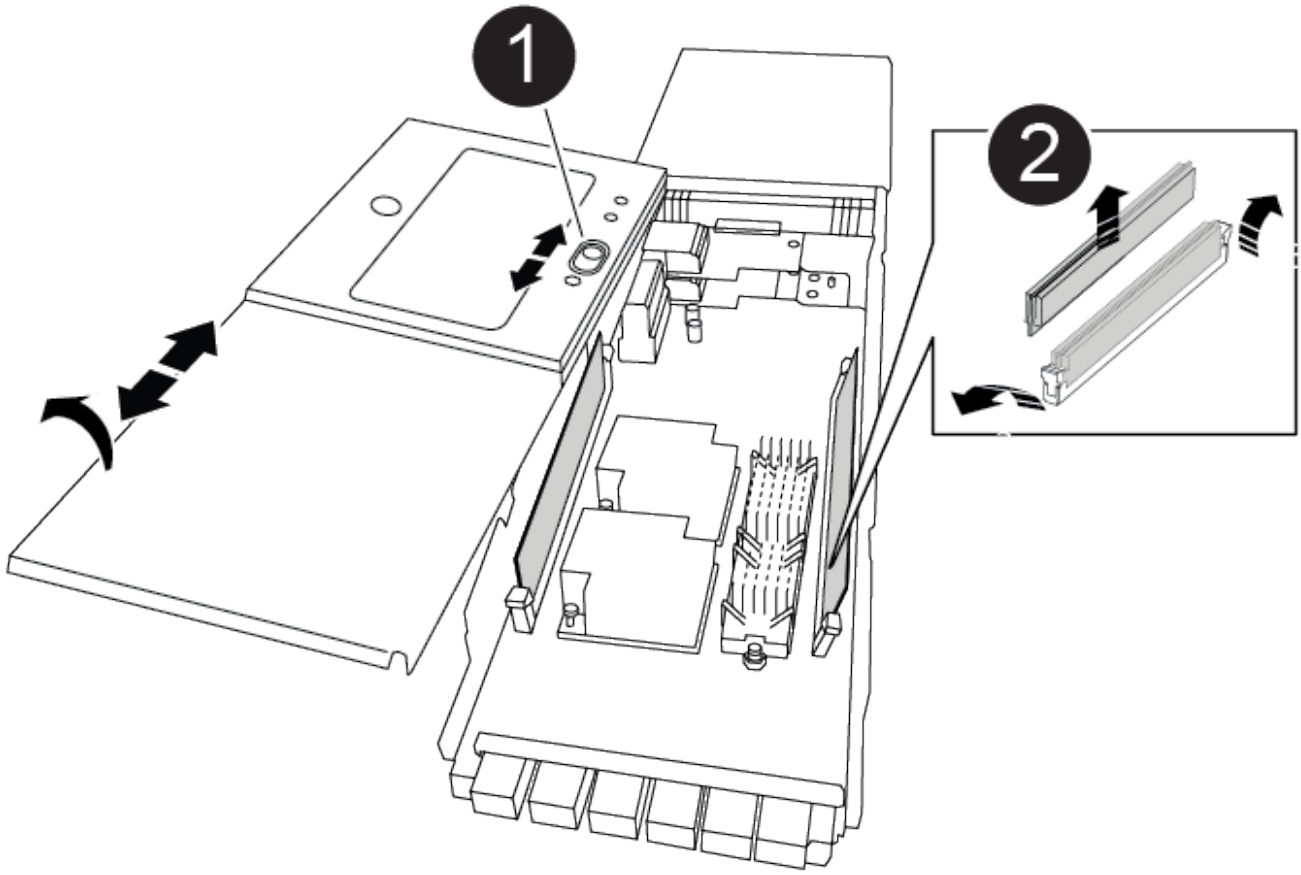

1	Cover locking button
2	DIMM and DIMM ejector tabs

4. Remove the DIMMs, one at a time, from the old NVRAM module and install them in the replacement NVRAM module.
5. Close the cover on the module.
6. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered cam latch begins to engage with the I/O cam pin, and then push the cam latch all the way up to lock the module in place.

Step 3: Replace a NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, you must remove the NVRAM module, open the module, and then replace the target DIMM.

1. If you are not already grounded, properly ground yourself.
2. Remove the target NVRAM module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

b. Rotate the cam latch down until it is in a horizontal position.

   The NVRAM module disengages from the chassis and moves out a few inches.

c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

   Animation - Replace NVRAM DIMM

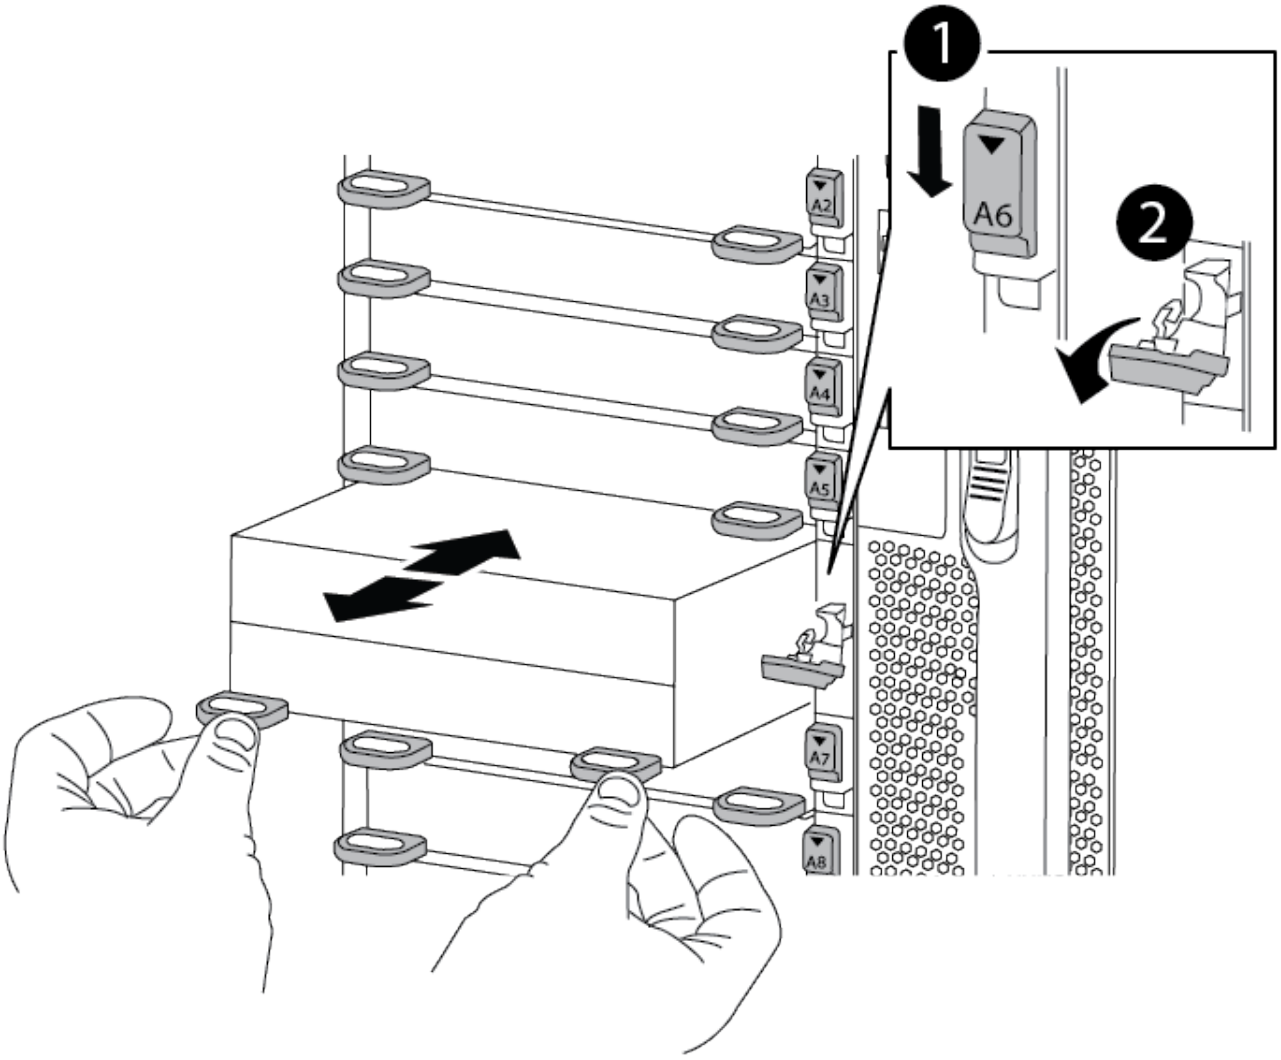

1	Lettered and numbered cam latch
2	cam latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

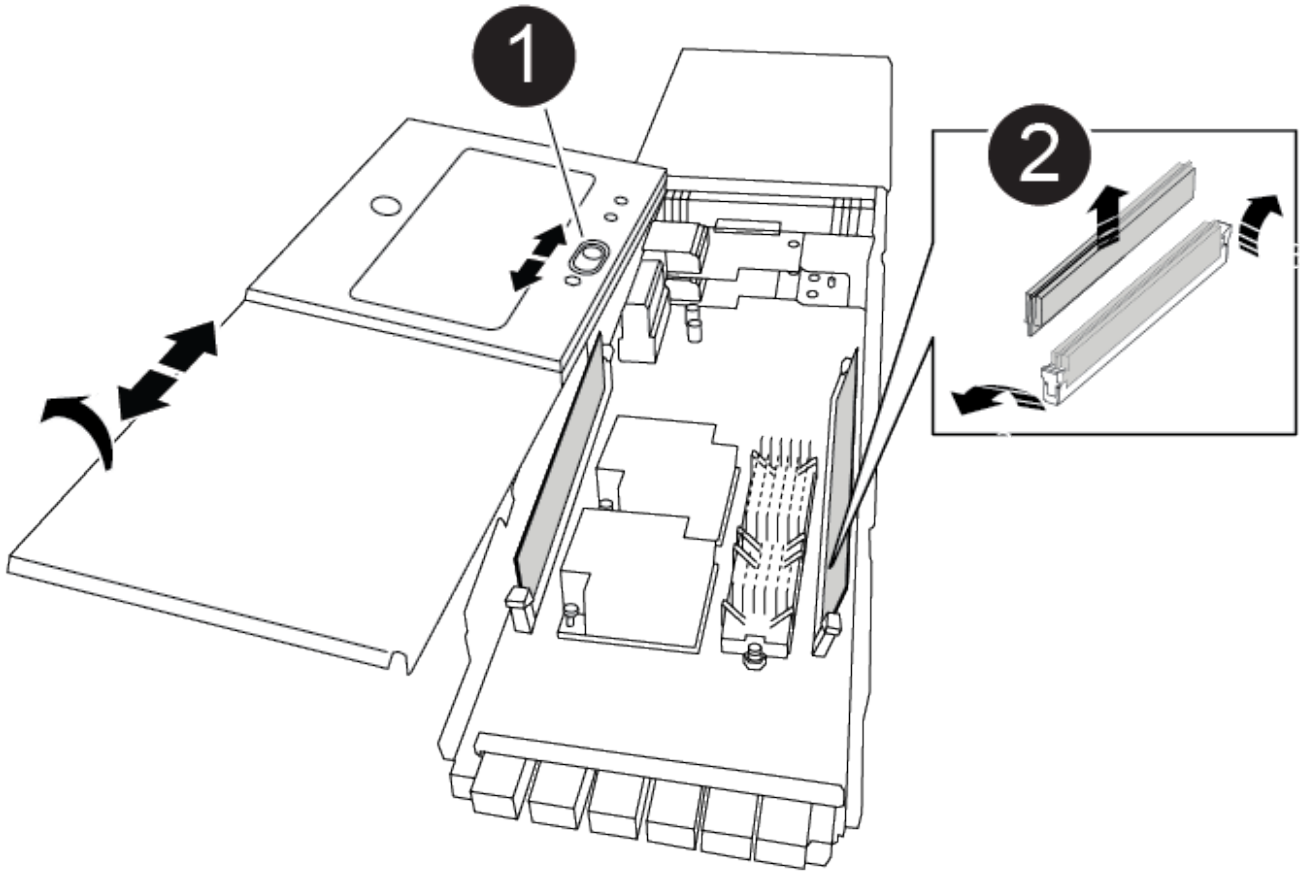

1	Cover locking button
2	DIMM and DIMM ejector tabs

4. Locate the DIMM to be replaced inside the NVRAM module, and then remove it by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
5. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
6. Close the cover on the module.
7. Install the NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered cam latch begins to engage with the I/O cam pin, and then push the cam latch all the way up to lock the module in place.

Step 4: Reboot the controller

After you replace the FRU, you must reboot the controller module.

1. To boot ONTAP from the LOADER prompt, enter `bye`.

Step 5: Reassign disks

You must confirm the system ID change when you boot the replacement controller and then verify that the change was implemented.

> Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the replacement controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the replacement controller, boot the controller and entering y if you are prompted to override the system ID due to a system ID mismatch.
3. Wait until the Waiting for giveback… message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1:> storage failover show
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The replacement controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

5. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the replacement controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 151759706:

```
node1:> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        151759706  151759706  -
151759706 Pool0
1.0.1  aggr0_1  node1 node1           151759706  151759706  -
151759706 Pool0
.
.
.
```

6. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

7. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The replacement controller is the current owner of the disks on the disaster site.

     See Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration for more information.

8. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

9. Verify that the expected volumes are present for each controller: `vol show -node node-name`
10. If storage encryption is enabled, you must restore functionality.
11. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - ASA A900

Swapping out a power supply involves turning off, disconnecting, and removing the power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

About this task

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- There are four power supplies in the system.
- Power supplies are auto-ranging.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.

3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
4. Press and hold the terra cotta button on the power supply handle, and then pull the power supply out of the chassis.

   CAUTION:

   When removing a power supply, always use two hands to support its weight.

   Animation - Remove/install PSU

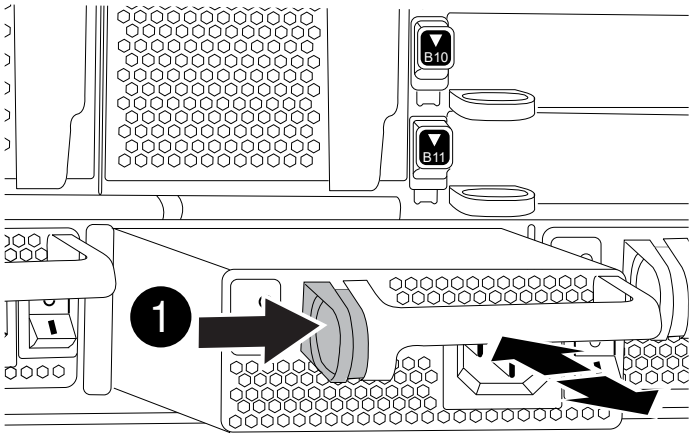

1	Locking button

5. Make sure that the on/off switch of the new power supply is in the Off position.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply.
   b. Secure the power cable to the power supply using the power cable retainer.

      Once power is restored to the power supply, the status LED should be green.

8. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

The green power LED lights when the PSU is fully inserted into the chassis and the amber attention LED flashes initially, but turns off after a few moments.

9. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA A900

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove the controller

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

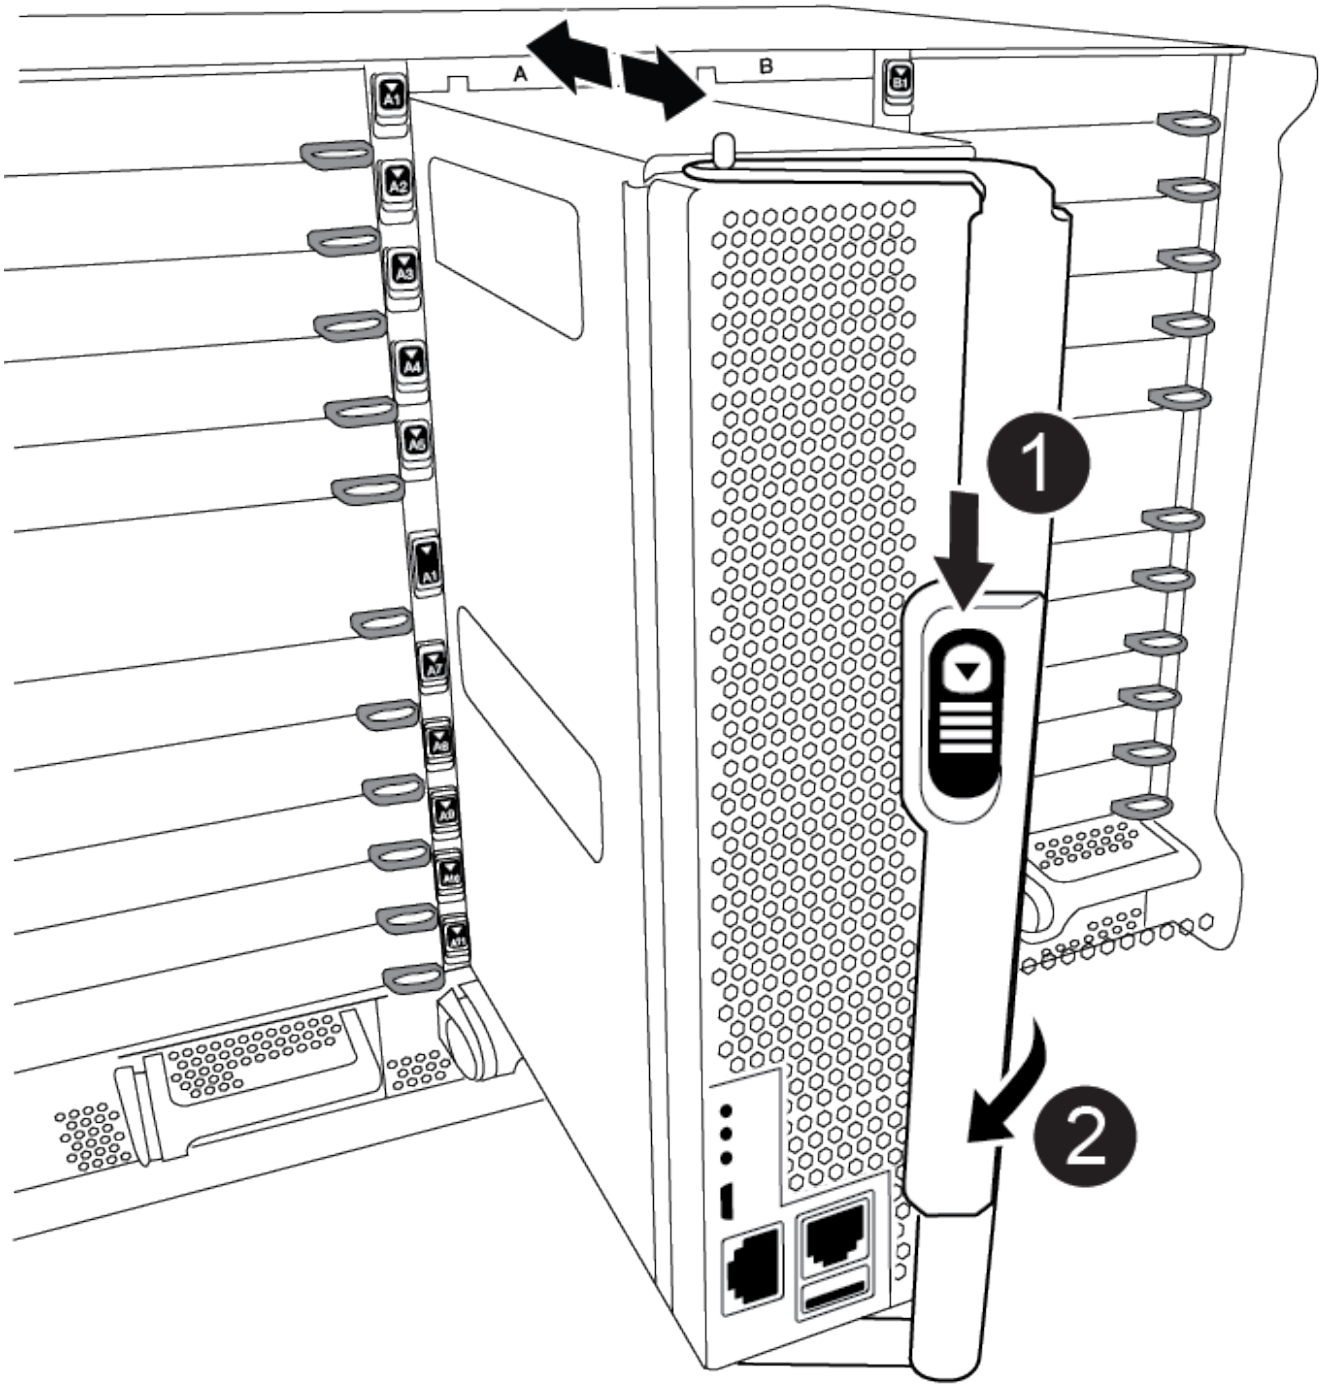

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

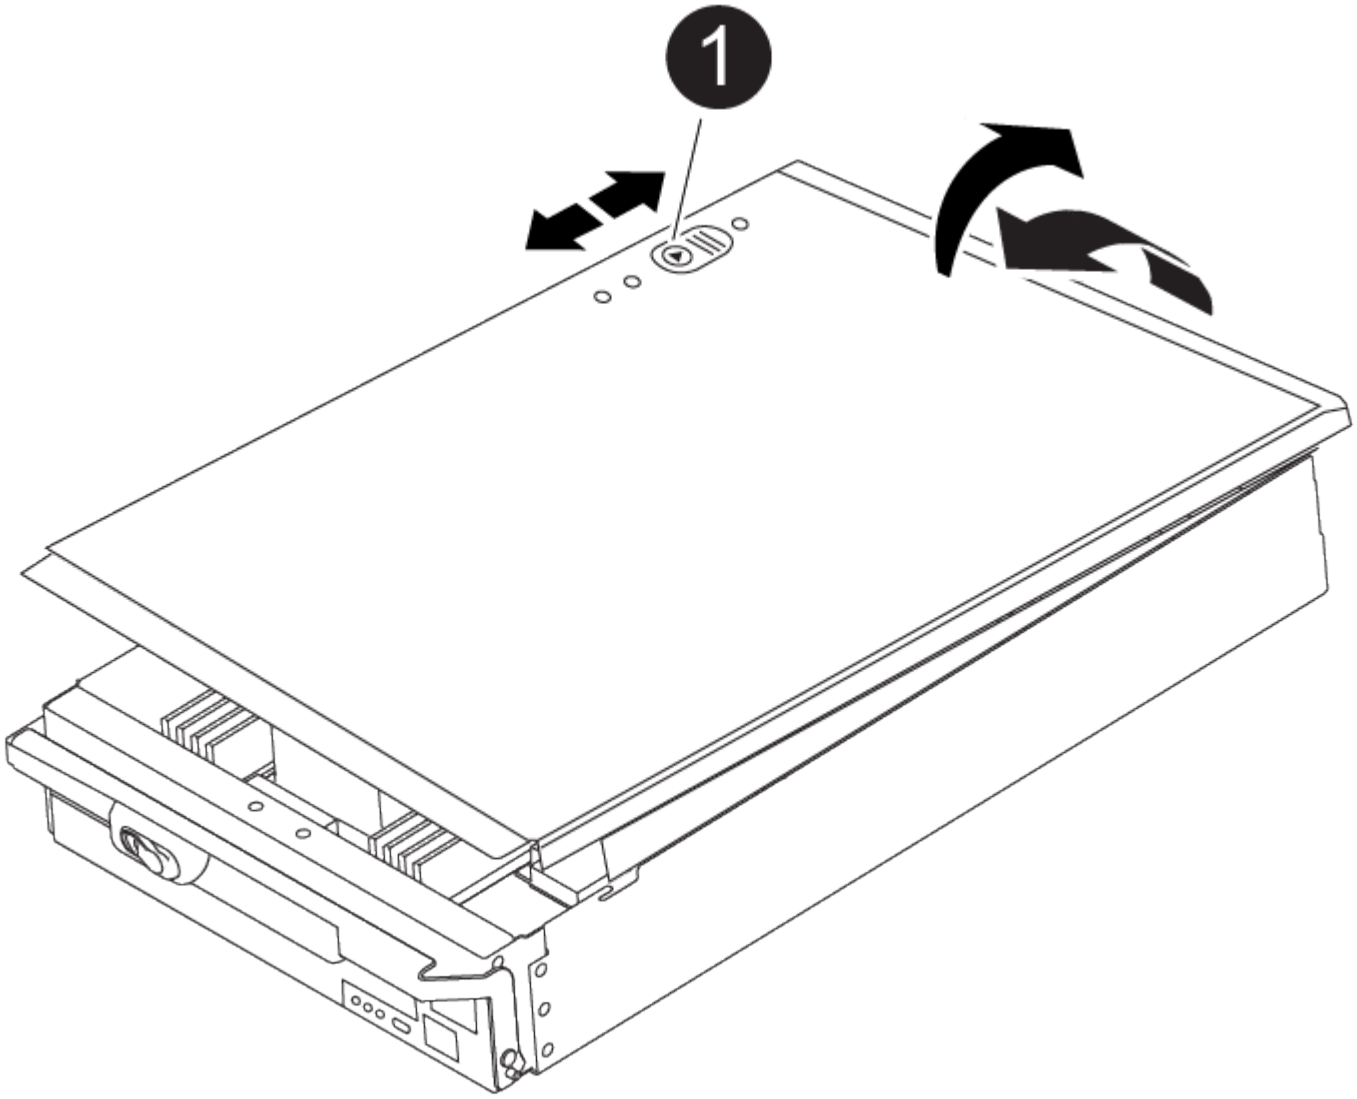

1	Controller module cover locking button

Step 3: Replace the RTC battery

To replace the RTC battery, you must locate the failed battery in the controller module, remove it from the holder, and then install the replacement battery in the holder.

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

   Animation - Replace RTC battery

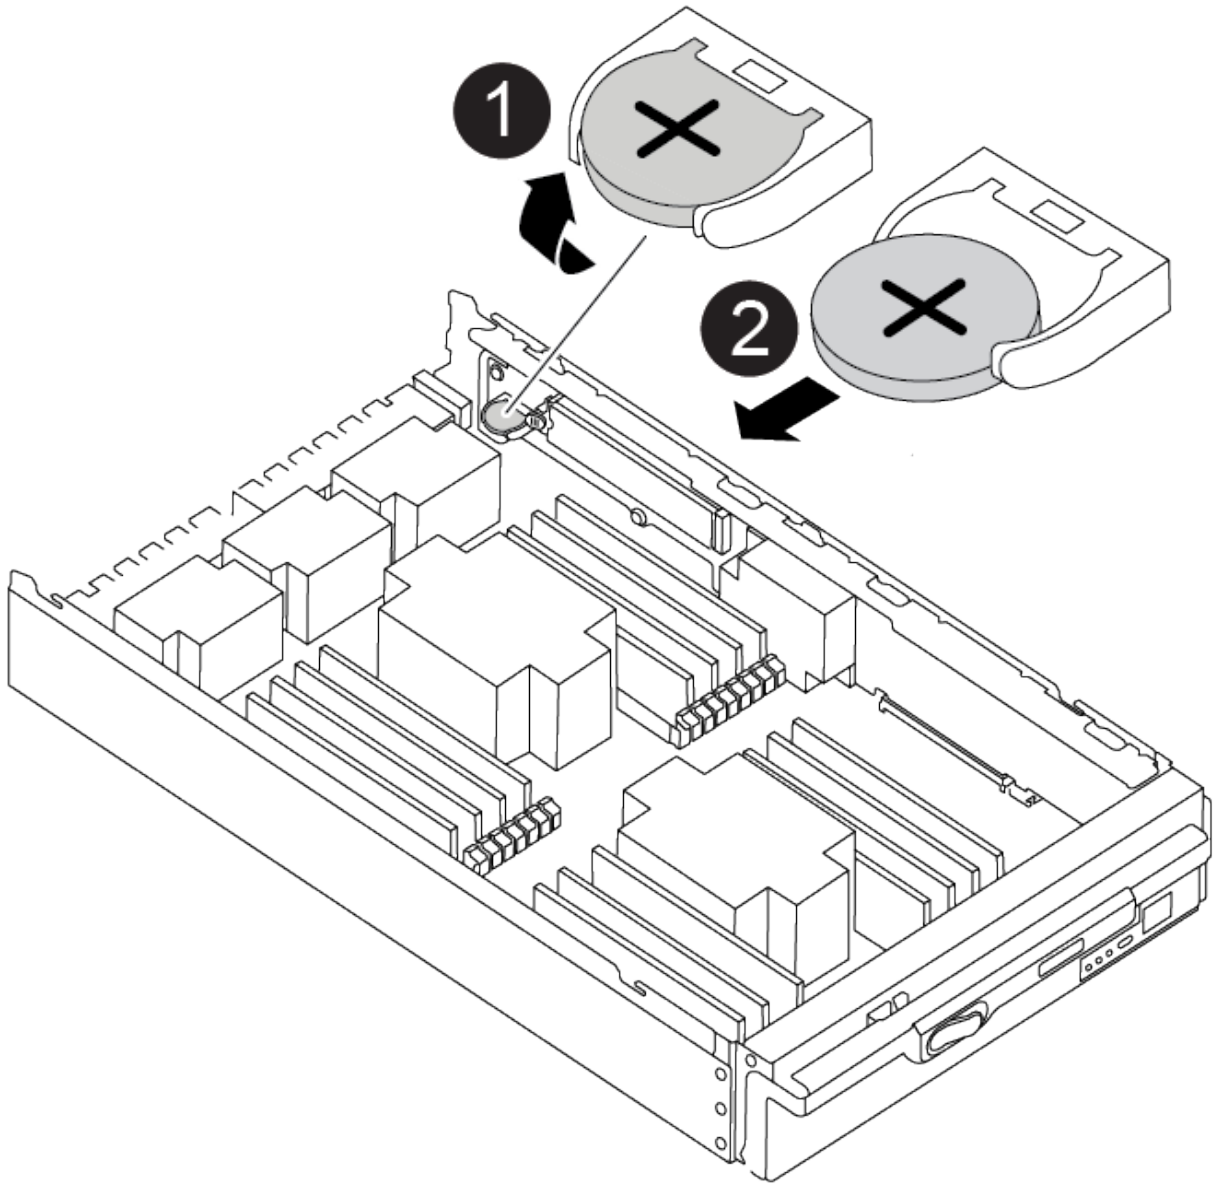

1	RTC battery
2	RTC battery housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.

7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
8. Reinstall the controller module cover.

Step 4: Reinstall the controller module and set time/date

After you replace the RTC battery, you must reinstall the controller module. If the RTC battery has been left out of the controller module for more than 10 minutes, you may have to reset the time and date.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.

If your system stops at the boot menu, select the option for "Reboot node" and respond y when prompted,
then boot to LOADER by pressing `Ctrl-C`.

1. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
2. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

3. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA C-Series systems

ASA C250 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick steps - ASA C250

The Installation and Setup instructions give graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

The ASA A250 and ASA C250 use the same installation procedure as the AFF A250 system.

AFF A250 Installation and Setup Instructions

Videos - ASA C250

The following video shows how to install and cable your new system.

Animation - Installation and Setup of an AFF A250

The ASA C250 uses the same installation procedure as the AFF A250 system.

Detailed steps - ASA C250

This section gives detailed step-by-step instructions for installing an ASA C250 system.

Step 1: Prepare for installation

To install your AFF A250 system, you need to create an account and register the system. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Customers with specific power requirements must check HWU for their configuration options.

Before you begin

- Make sure you have access to the NetApp Hardware Universe (HWU) for information about site requirements.
- Make sure you have access to the Release Notes for your version of ONTAP for more information about this system.
- You need to provide the following at your site:
  - Rack space for the storage system
  - Phillips #2 screwdriver
  - Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register your system.
4. Download and install NetApp Downloads: Config Advisor on your laptop.
5. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
25 GbE cable	X66240A-05 (112-00595), 0.5m; X66240-2 (112-00573), 2m		Cluster interconnect network
	X66240A-2 (112-00598), 2m; X66240A-5 (112-00600), 5m		Data
100 GbE cable	X66211-2 (112-00574), 2m; X66211-5 (112-00576), 5m		Storage
RJ-45 (order dependent)	Not applicable		Management network (BMC and wrench port) and Ethernet data (e0a and e0b)
Fibre Channel	X66250-2 (112-00342) 2m; X66250-5 (112-00344) 5m; X66250-15 (112-00346) 15m; X66250-30 (112-00347) 30m		
Micro-USB console cable	Not applicable		Console connection during software setup
Power cables	Not applicable		Powering up the system

6. Review the ONTAP Configuration Guide and collect the required information listed in that guide.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

> You need to be aware of the safety concerns associated with the weight of the system.

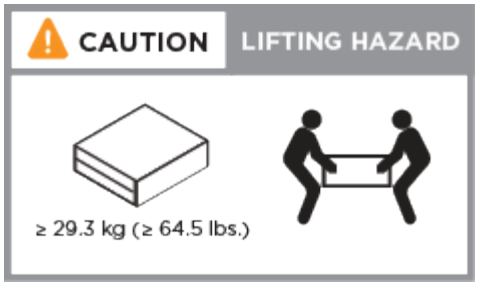

3. Identify and manage cables because this system does not have a cable management device.
4. Place the bezel on the front of the system.

Step 3: Cable controllers to cluster

Cable the controllers to a cluster by using the two-node switchless cluster method or by using the cluster interconnect network method.

Option 1: Two-node switchless cluster

The management, Fibre Channel, and data or host network ports on the controller modules are connected to switches. The cluster interconnect ports are cabled on both controller modules.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

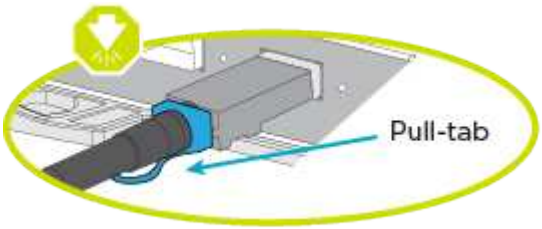

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the tabulated steps to complete the cabling between the controllers and the switches. Perform the steps on each controller.

Animation - Cable two-node switchless cluster

Steps

1. Use the the 25GbE cluster interconnect cable to connect the cluster interconnect ports e0c to e0c and e0d to e0d.


:

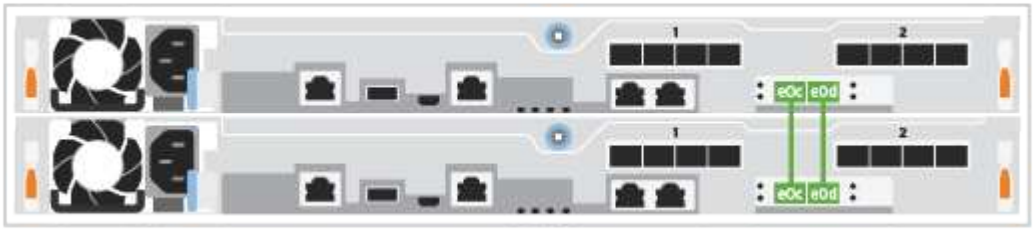

2. Cable the wrench ports to the management network switches with the RJ45 cables.

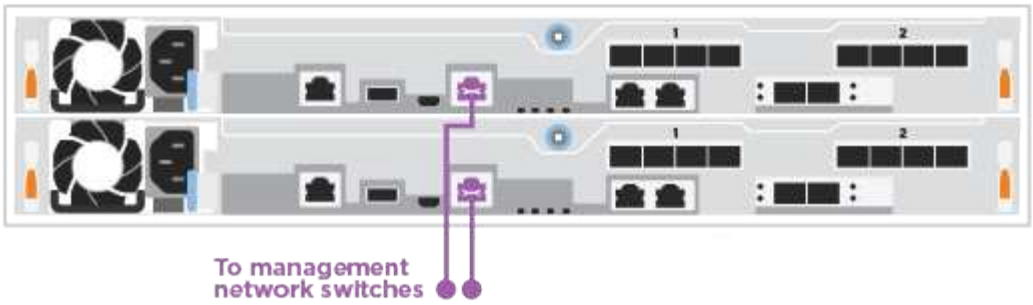

DO NOT plug in the power cords at this point.

Option 2: Switched cluster

All ports on the controllers are connected to switches; cluster interconnect, management, Fibre Channel, and data or host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

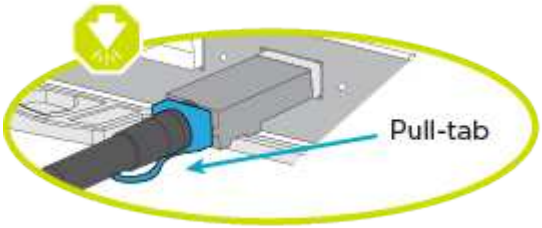

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the steps to complete the cabling between the controllers and the switches. Perform the steps on each controller.

Animation - Cable switched cluster

Steps

1. Cable the cluster interconnect ports e0c and e0d to the 25 GbE cluster interconnect switches.

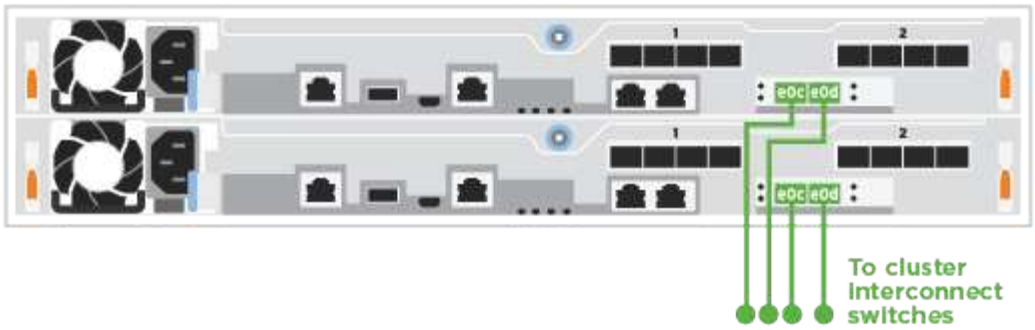

2. Cable the wrench ports to the management network switches with the RJ45 cables.

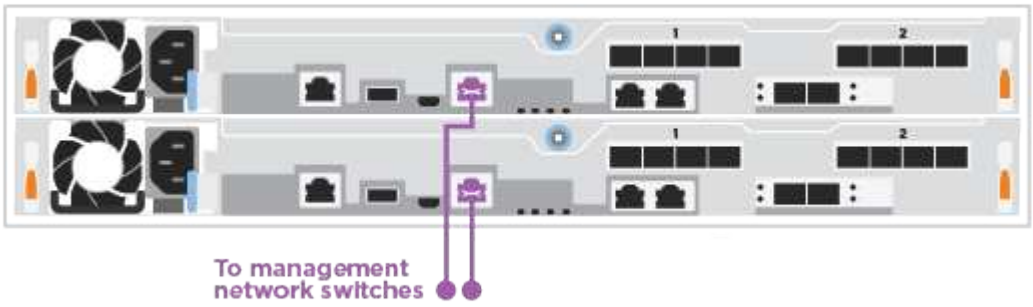

Step 4: Cable to host network or storage (Optional)

You have configuration-dependent optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cabling to a host network and storage.

NetApp Hardware Universe slot priority for host network cards (Fibre Channel or 25GbE) is slot 2. However, if you have both cards, the Fibre Channel card goes in slot 2 and the 25GbE card goes in slot 1 (as shown in the options below). If you have an external shelf, the storage card goes in slot 1, the only supported slot for shelves.

Option 1: Cable to Fibre Channel host network

Fibre Channel ports on the controllers are connected to Fibre Channel host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

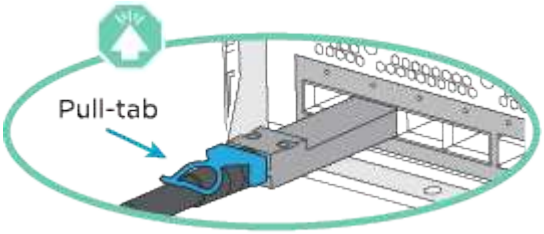

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again..

About this task

Perform the following step on each controller module.

Steps

1. Cable ports 2a through 2d to the FC host switches.

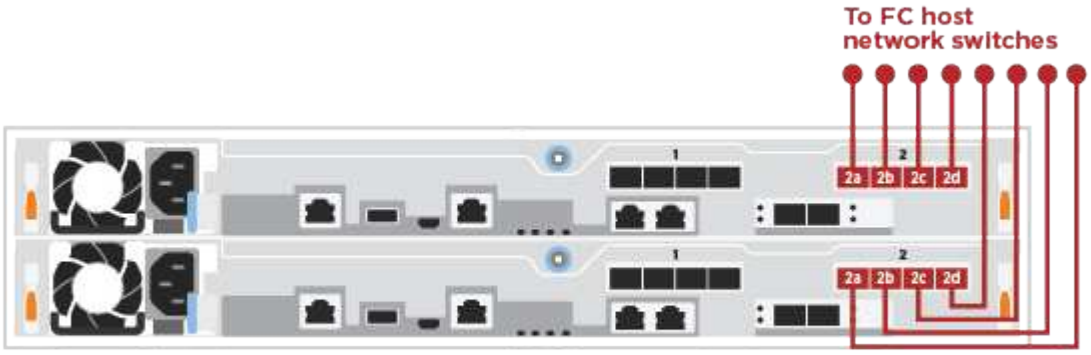

Option 2: Cable to 25GbE data or host network

25GbE ports on the controllers are connected to 25GbE data or host network switches.

Before you begin

- Contact your network administrator for information about connecting the system to the switches.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

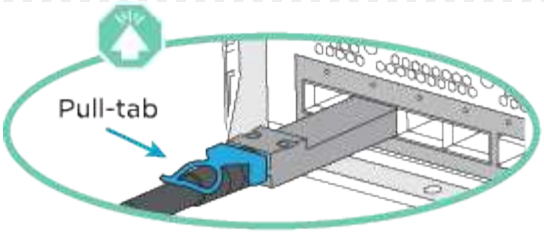

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Perform the following step on each controller module.

Steps

1. Cable ports e4a through e4d to the 10GbE host network switches.

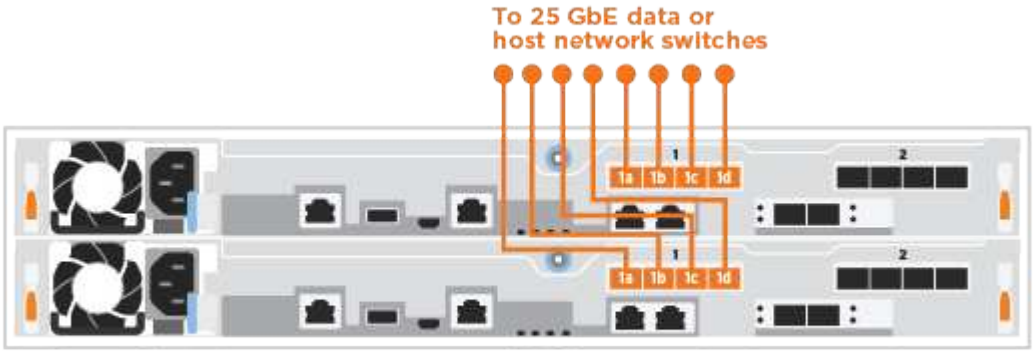

Option 3: Cable controllers to single drive shelf

Cable each controller to the NSM modules on the NS224 drive shelf.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

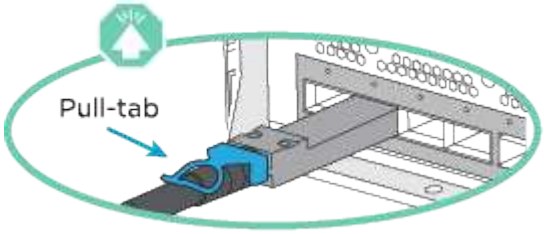

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

About this task

Use the animation or the tabulated steps to complete the cabling between the controllers and the single shelf. Perform the steps on each controller module.

Animation - Cable the controllers to a single NS224

Steps

1. Cable controller A to the shelf.

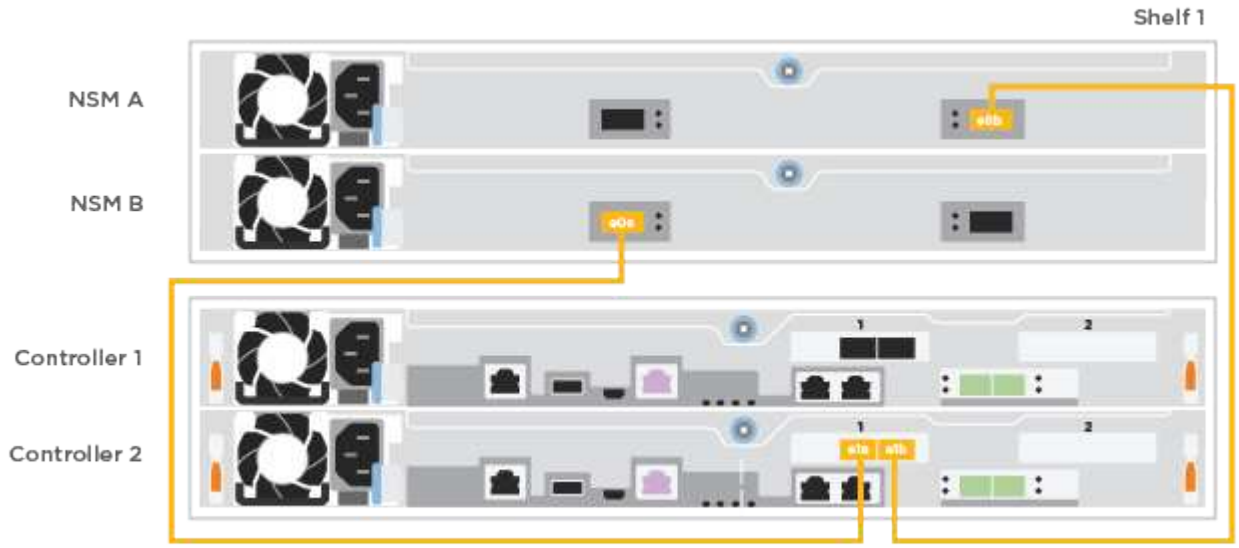

2. Cable controller B to the shelf.

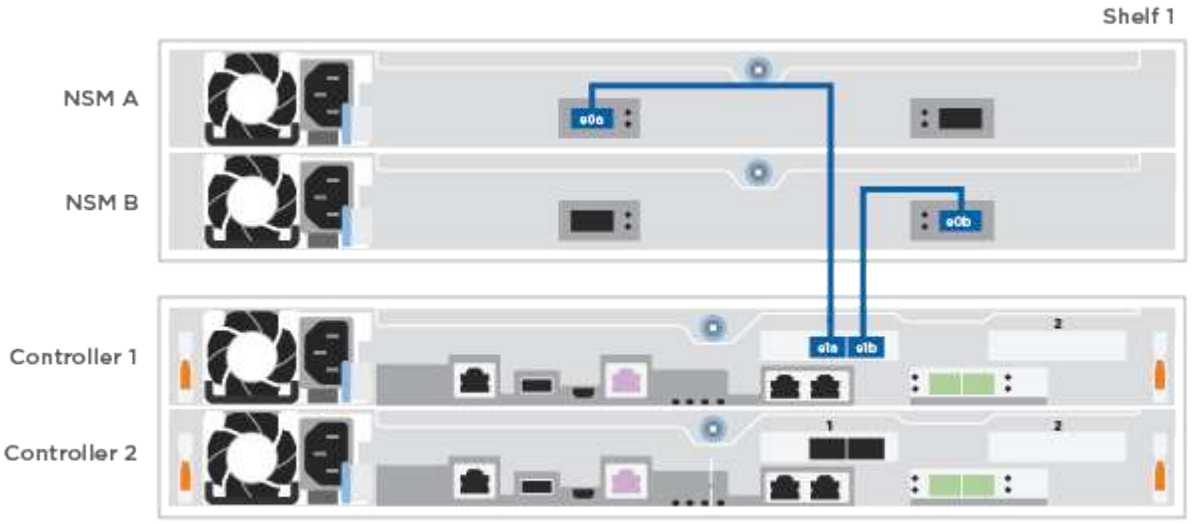

Step 5: Complete system setup

Complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: If network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

2. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

3. Use the animation to connect your laptop to the Management switch:

   Animation - Connect your laptop to the Management switch

4. Select an ONTAP icon listed to discover:

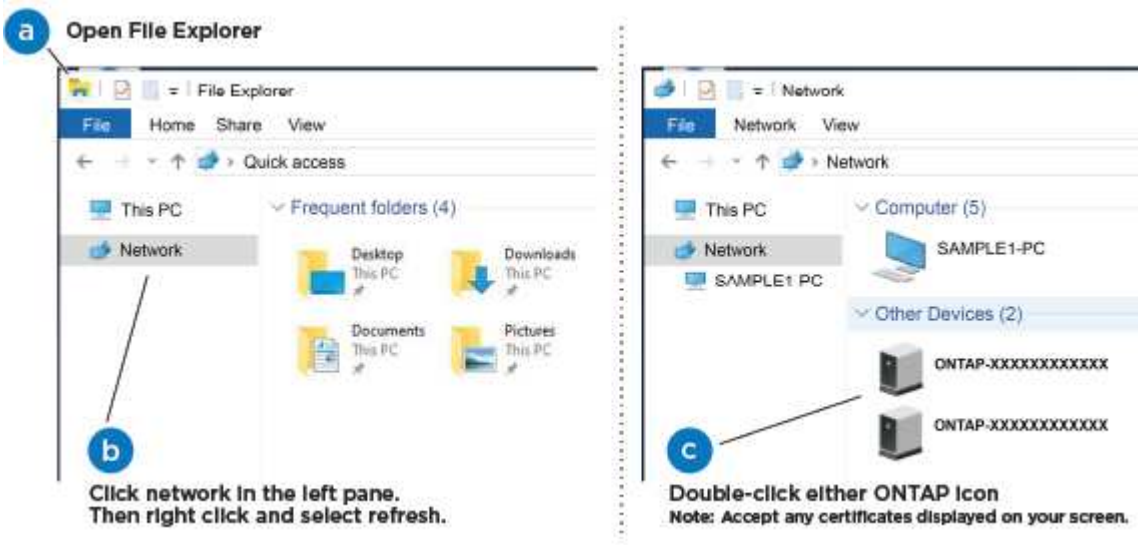

   a. Open File Explorer.

   b. Click **Network** in the left pane and right-click and select **refresh**.

   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

   System Manager opens.

5. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
6. Verify the health of your system by running Config Advisor.
7. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the laptop or console to the switch on the management subnet.

   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.
3. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

4. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the ONTAP Configuration Guide.

5. Verify the health of your system by running Config Advisor.
6. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain ASA C250 hardware

For the ASA C250 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

Fan

The fan cools the controller.

Mezzanine card

A Mezzanine card is a printed circuit board that plugs directly into another plug-in card.

NVEM battery

A battery is included with the controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real-time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - ASA C250

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.
- You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

About this task

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
  - The *impaired* node is the controller on which you are performing maintenance.
  - The *healthy* node is the HA partner of the impaired controller.

Check encryption key support and status - ASA C250

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select

one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller - ASA C250

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller displays…	Then…
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Option 2: Systems in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *impaired_node_name* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - ASA C250

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller module, you must first remove the controller module from the system, and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Unplug the I/O cables from the controller module.
5. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

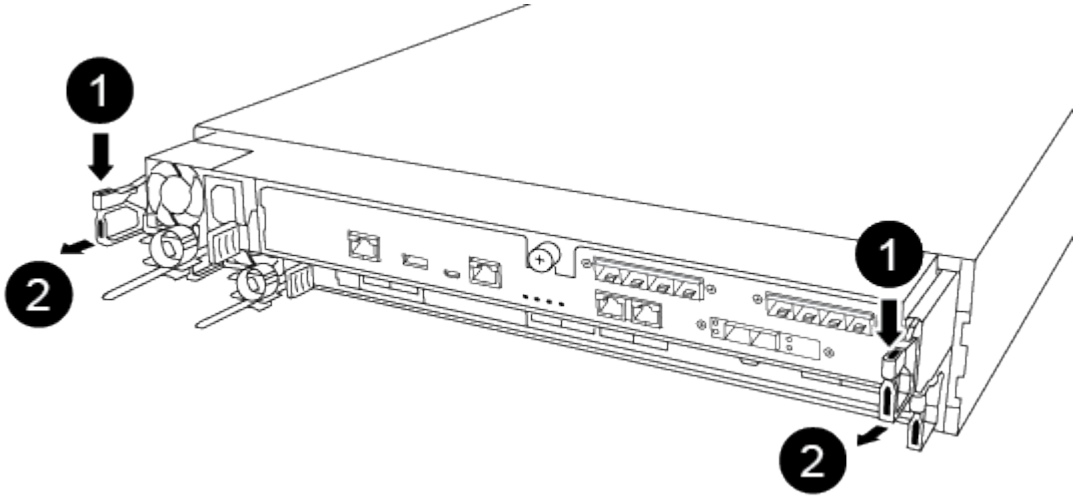

1	Lever
2	Latching mechanism

6. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
7. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

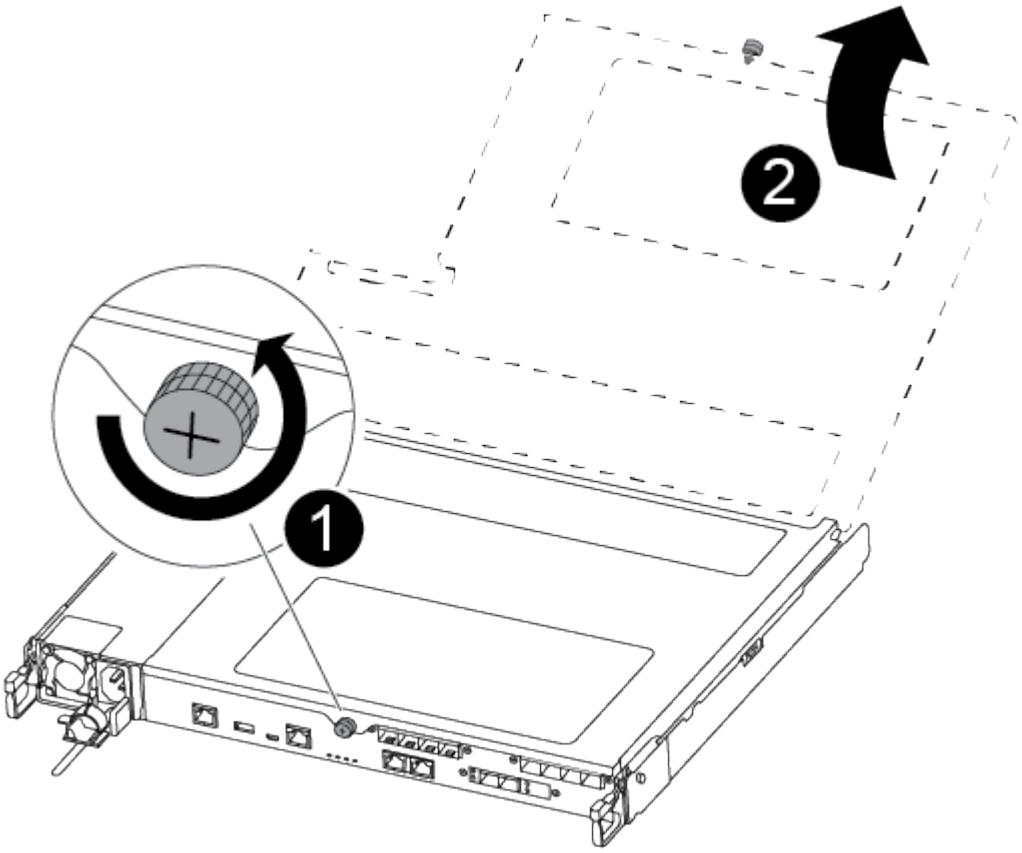

1	Thumbscrew
2	Controller module cover.

8. Lift out the air duct cover.

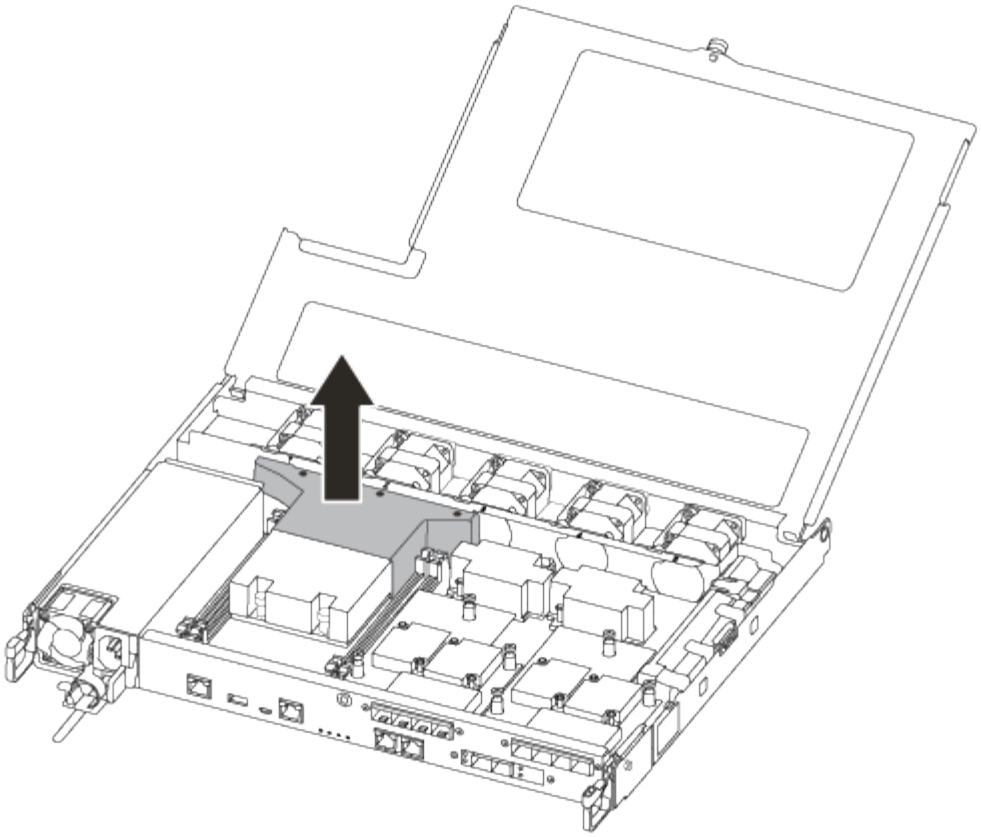

Step 2: Replace the boot media

You locate the failed boot media in the controller module by removing the air duct on the controller module before you can replace the boot media.

You need a #1 magnetic Phillips head screwdriver to remove the screw that holds the boot media in place. Due to the space constraints within the controller module, you should also have a magnet to transfer the screw on to so that you do not lose it.

You can use the following video or the tabulated steps to replace the boot media:

Animation - Replace the boot media

1. Locate and replace the impaired boot media from the controller module.

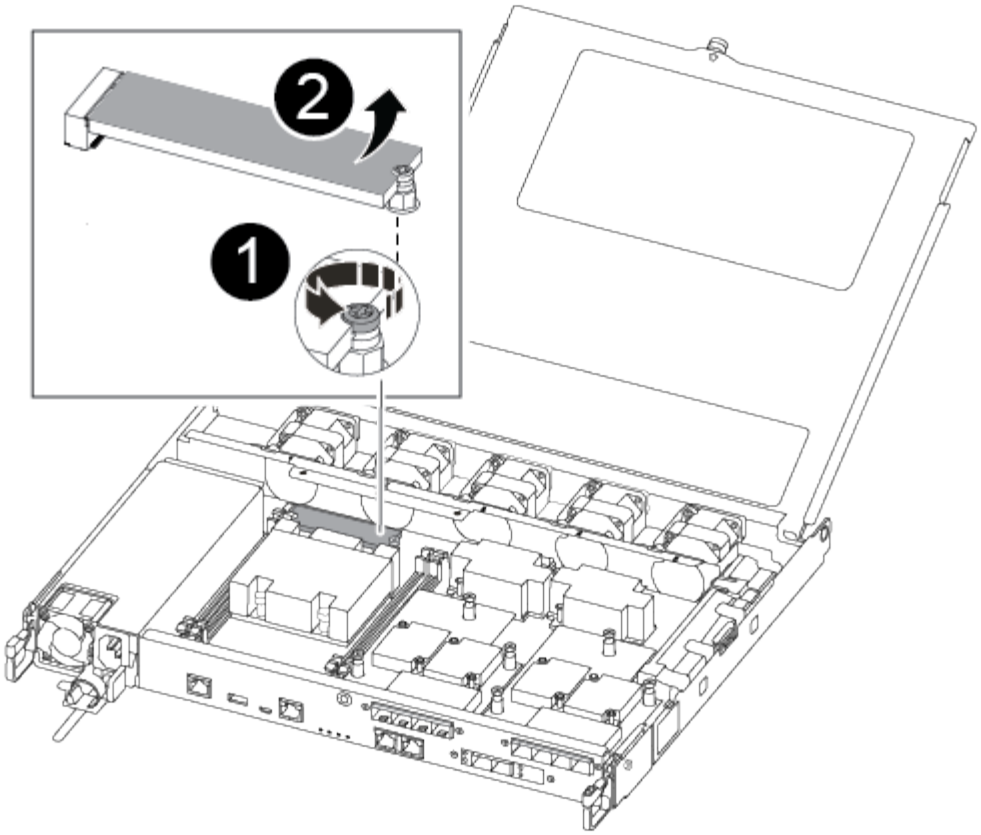

1	Remove the screw securing the boot media to the motherboard in the controller module.
2	Lift the boot media out of the controller module.

2. Using the #1 magnetic screwdriver, remove the screw from the impaired boot media, and set it aside safely on the magnet.
3. Gently lift the impaired boot media directly out of the socket and set it aside.
4. Remove the replacement boot media from the antistatic shipping bag and align it into place on the controller module.
5. Using the #1 magnetic screwdriver, insert and tighten the screw on the boot media.

> Do not apply force when tightening the screw on the boot media; you might crack it.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed is without a boot image so you need to transfer a boot image using a USB flash drive.

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download

button.

- If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
2. Download the service image to your work space on your laptop.
3. Unzip the service image.

   If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

   There are two folders in the unzipped service image file:

   - boot
   - efi

4. Copy the efi folder to the top directory on the USB flash drive.

   If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

   The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

5. Remove the USB flash drive from your laptop.
6. If you have not already done so, install the air duct.

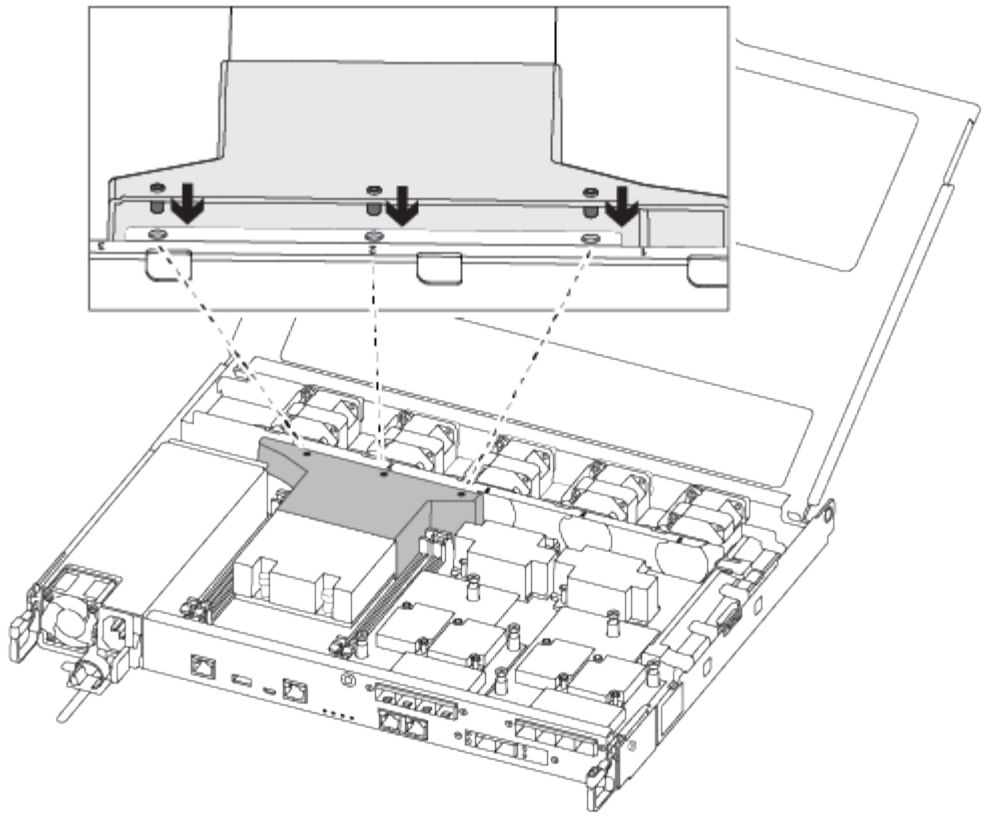

7. Close the controller module cover and tighten the thumbscrew.

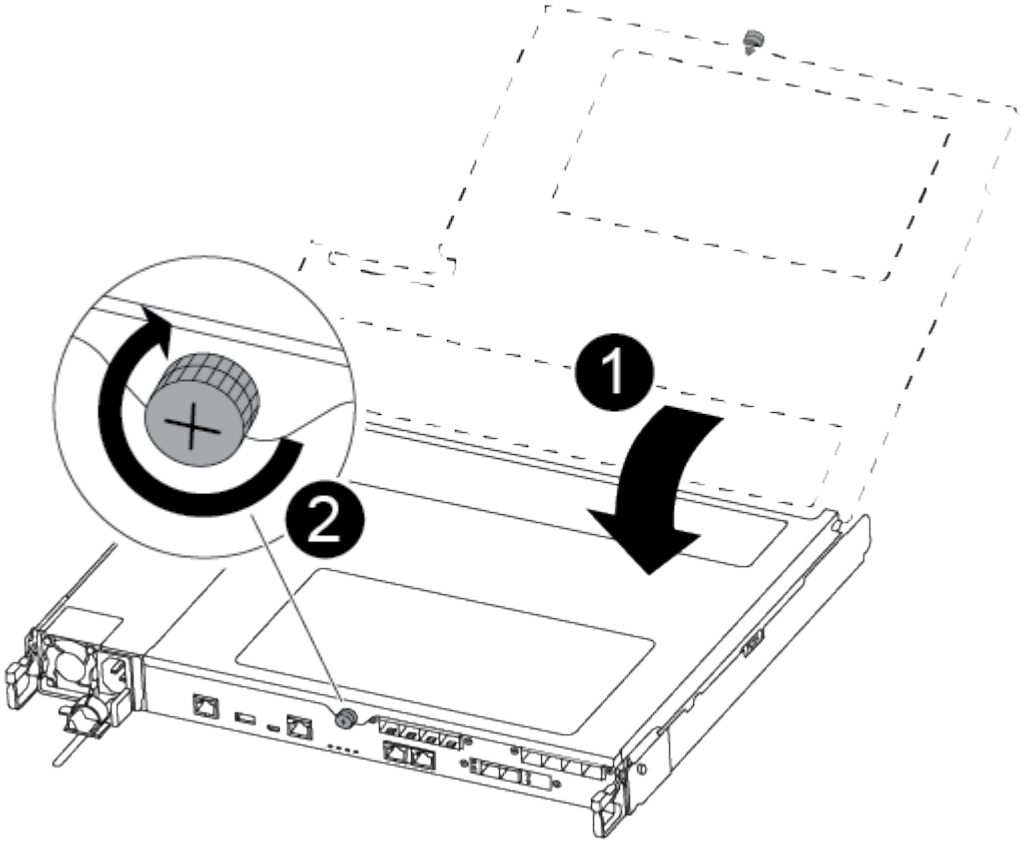

1	Controller module cover
2	Thumbscrew

8. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
9. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

10. Push the controller module all the way into the chassis:
11. Place your index fingers through the finger holes from the inside of the latching mechanism.
12. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
13. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

    The controller module should be fully inserted and flush with the edges of the chassis.

14. Reconnect the controller module I/O cables.
15. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

    The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

16. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

    If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

17. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

    The system begins to boot and stops at the LOADER prompt.

Boot the recovery image - ASA C250

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.

3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.

5. Give back the controller using the `storage failover giveback -fromnode local` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

   NOTE: If the process fails, contact NetApp Support.

Restore encryption - ASA C250

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
**********************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
**********************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - ASA C250

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - ASA C250

To replace the chassis, you must move the bezel, controller modules, and NVMe drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in its place.

About this task

- All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - ASA C250

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

```
Warning: Are you sure you want to halt node <node_name>? {y|n}:
```

10. Wait for each controller to halt and display the LOADER prompt.

Replace hardware - ASA C250

To replace the chassis, you move the power supplies, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis from with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

Use the following video or the tabulated steps to replace the chassis; it assumes the removal and replacement of the bezel:

Animation - Replace the chassis

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

> If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

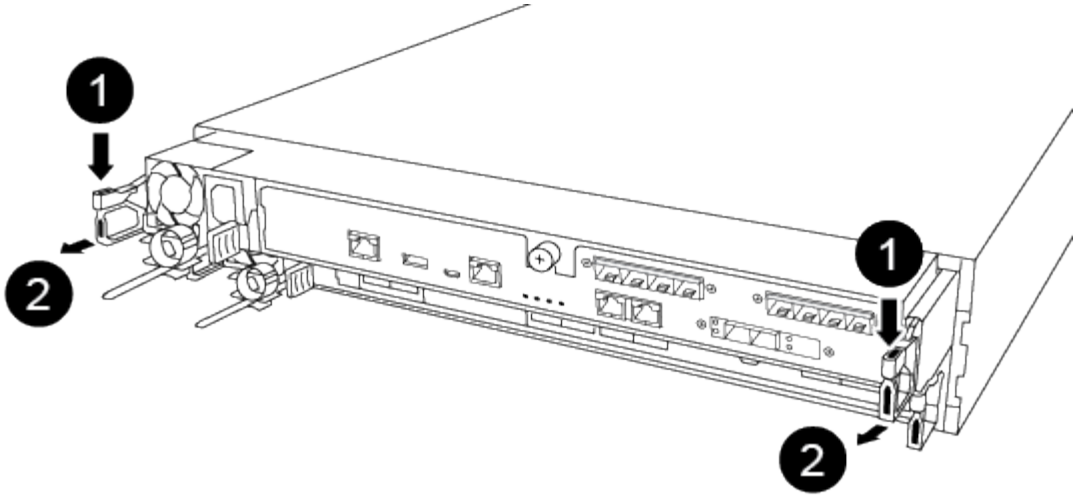

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.
5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up

and against the drive holder.

Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot the system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Insert the controller module into the chassis:

   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

   The controller module should be fully inserted and flush with the edges of the chassis.

   a. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect

the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - ASA C250

You must verify the HA state of the chassis, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement- ASA C250

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct

recovery procedure to determine whether you should use this procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller module - ASA C250

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - ASA C250

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

Use the following video or the tabulated steps to replace a controller module:

Animation - Replace a controller module

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

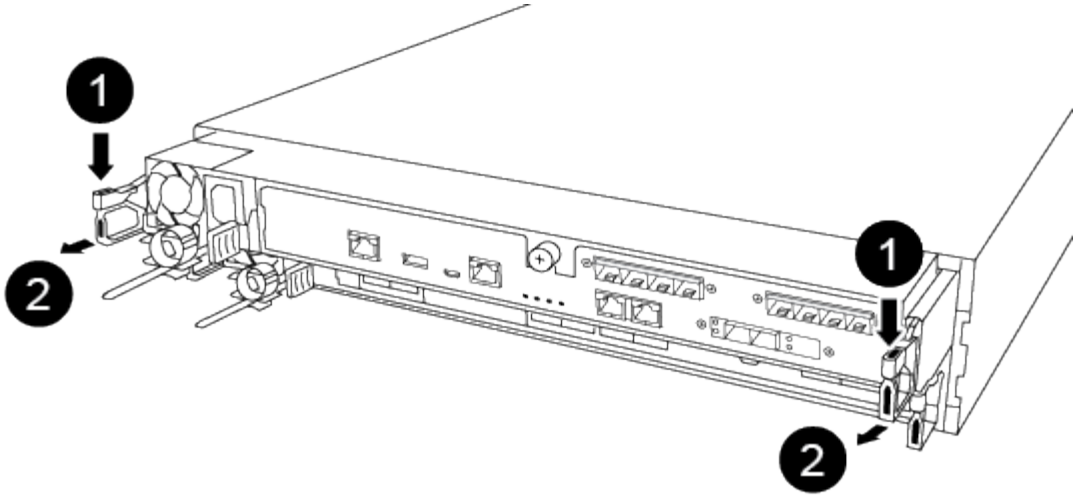

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

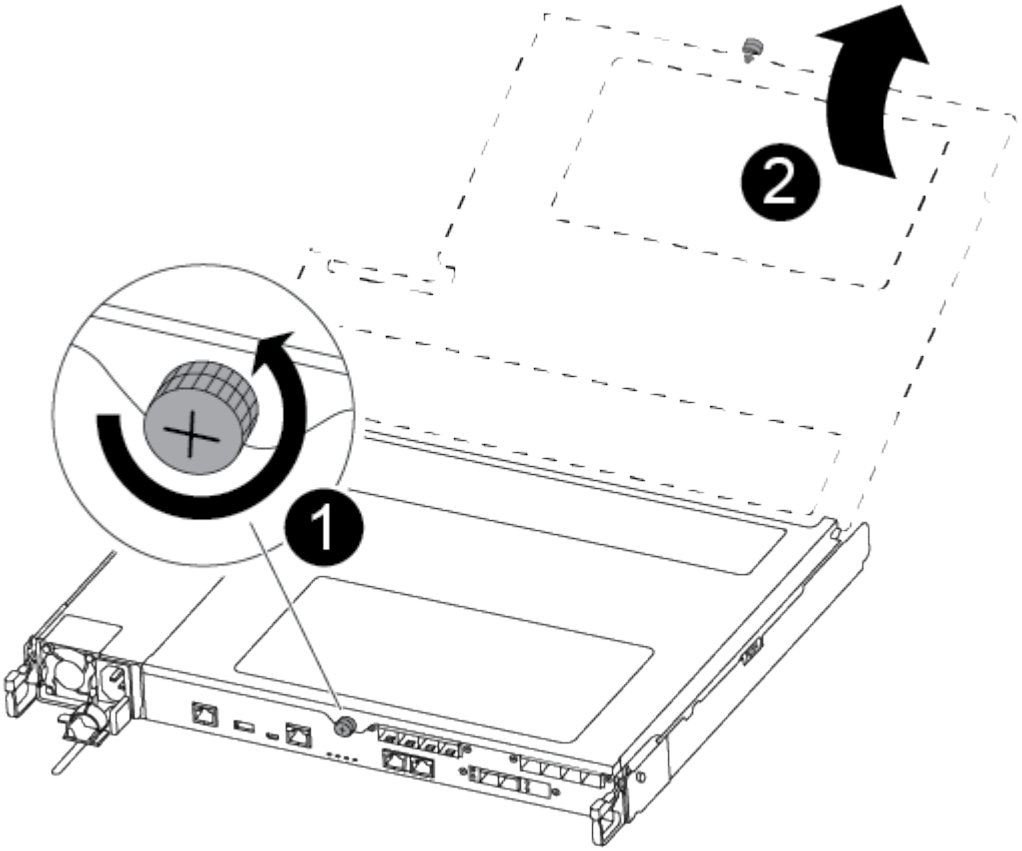

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

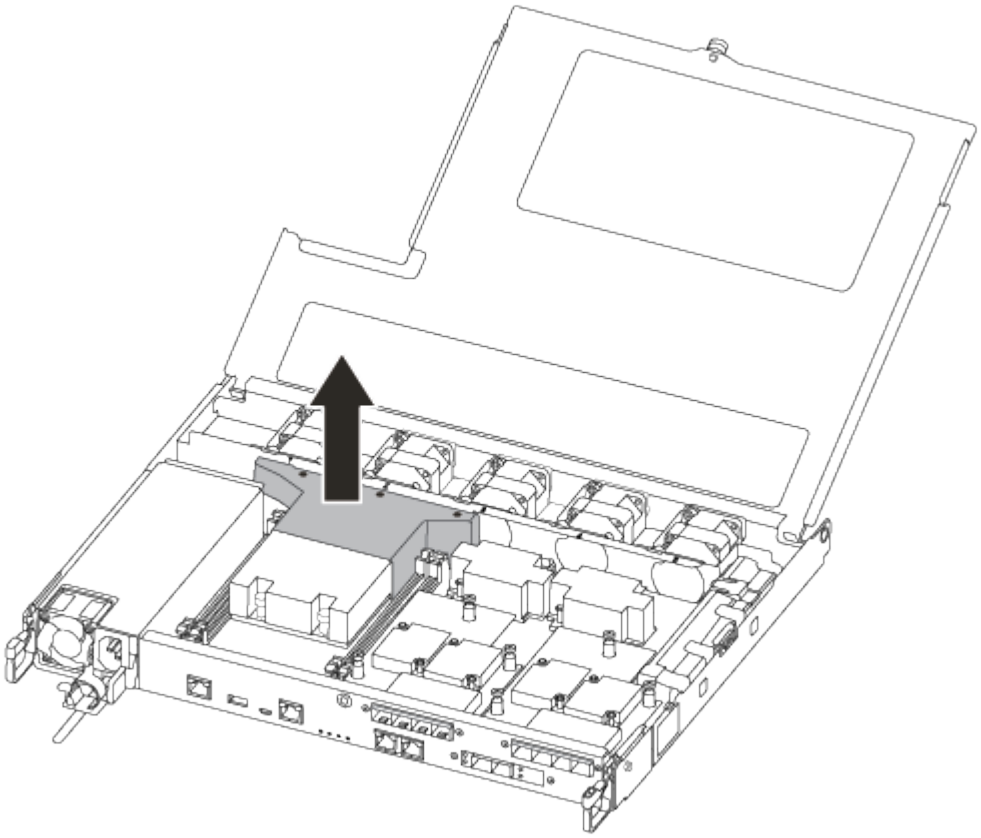

Step 2: Move the power supply

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

1. Disconnect the power supply.
2. Open the power cable retainer, and then unplug the power cable from the power supply.
3. Unplug the power cable from the power source.
4. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

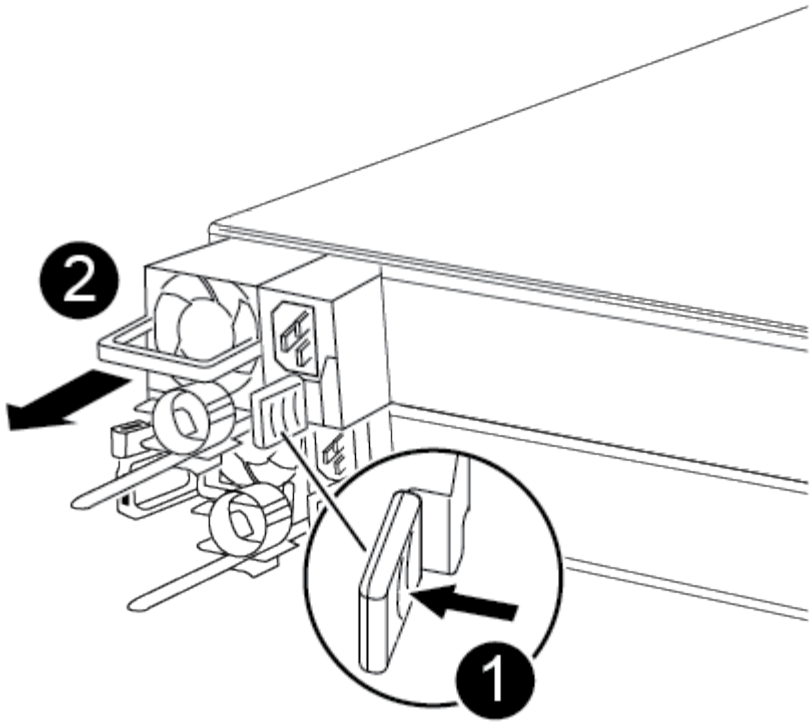

1	Blue power supply locking tab
2	Power supply

5. Move the power supply to the new controller module, and then install it.
6. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

You must move the fans from the impaired controller module to the replacement module when replacing a failed controller module.

1. Remove the fan module by pinching the side of the fan module, and then lifting the fan module straight out of the controller module.

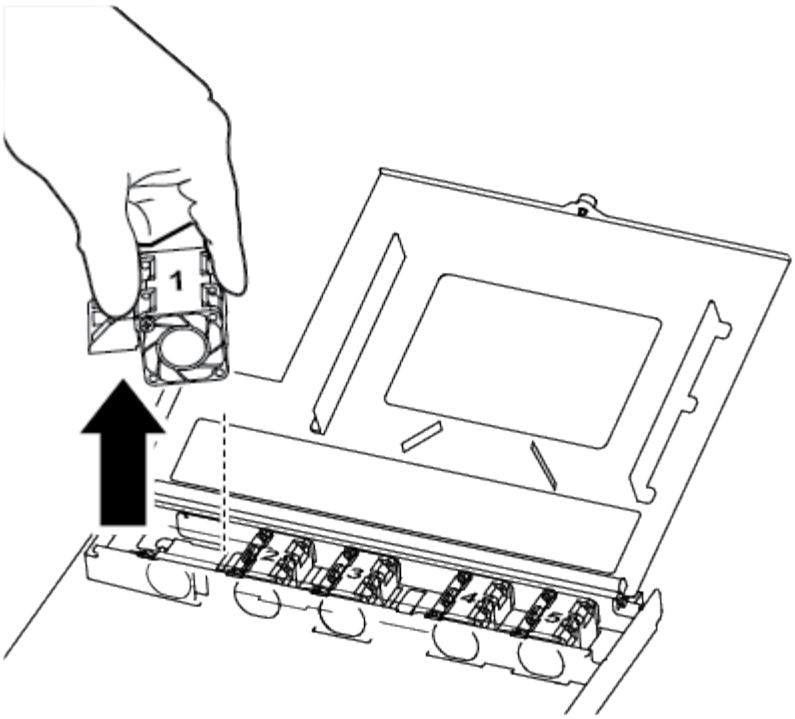

1	Fan module

2. Move the fan module to the replacement controller module, and align the edges of the fan module with the opening in the controller module, and then slide the fan module in.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the boot media

You must move the boot media device from the impaired controller module to the replacement controller module.

You need a #1 magnetic Phillips head screwdriver to remove the screw that holds the boot media in place. Due to the space constraints within the controller module, you should also have a magnet to transfer the screw on to so that you do not lose it.

1. Locate and move the boot media from the impaired controller module to the replacement controller module.

   The boot media is located under the air duct cover you removed earlier in this procedure.

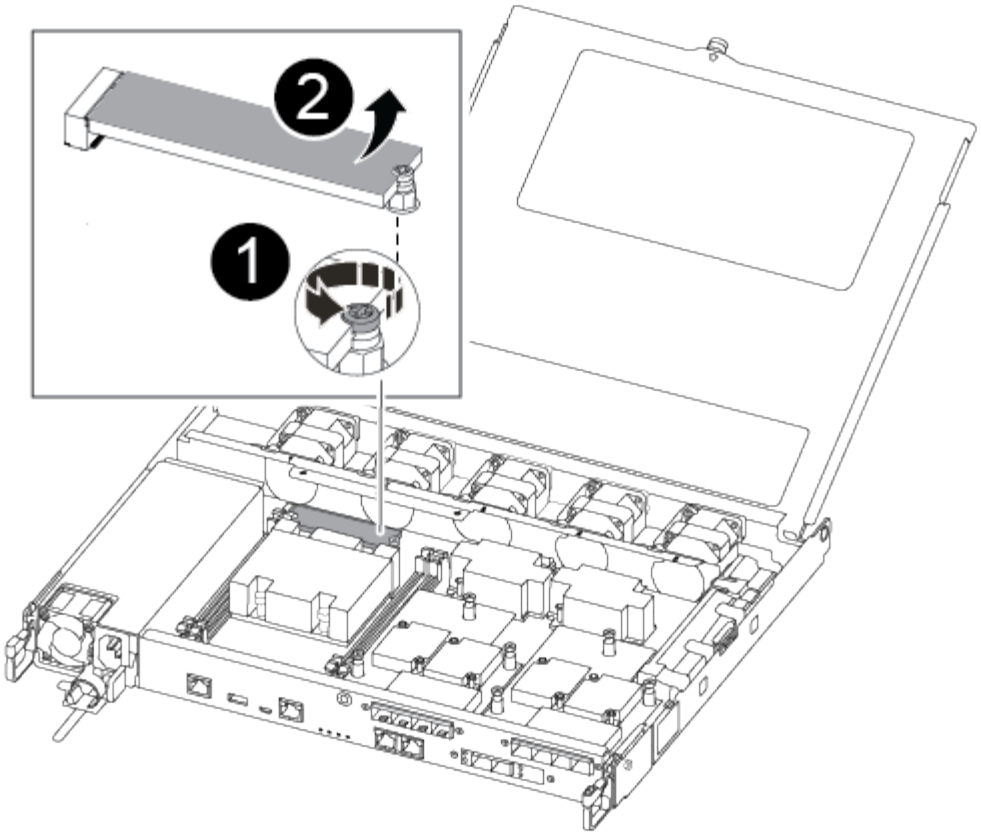

1	Remove the screw securing the boot media to the motherboard in the impaired controller module.
2	Lift the boot media out of the impaired controller module.

2. Using the #1 magnetic screwdriver, remove the screw from the boot media, and set it aside safely on the magnet.
3. Gently lift the boot media directly out of the socket and align it into place in the replacement controller module.
4. Using the #1 magnetic screwdriver, insert and tighten the screw on the boot media.

> Do not apply force when tightening the screw on the boot media; you might crack it.

Step 5: Move the DIMMs

To move the DIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

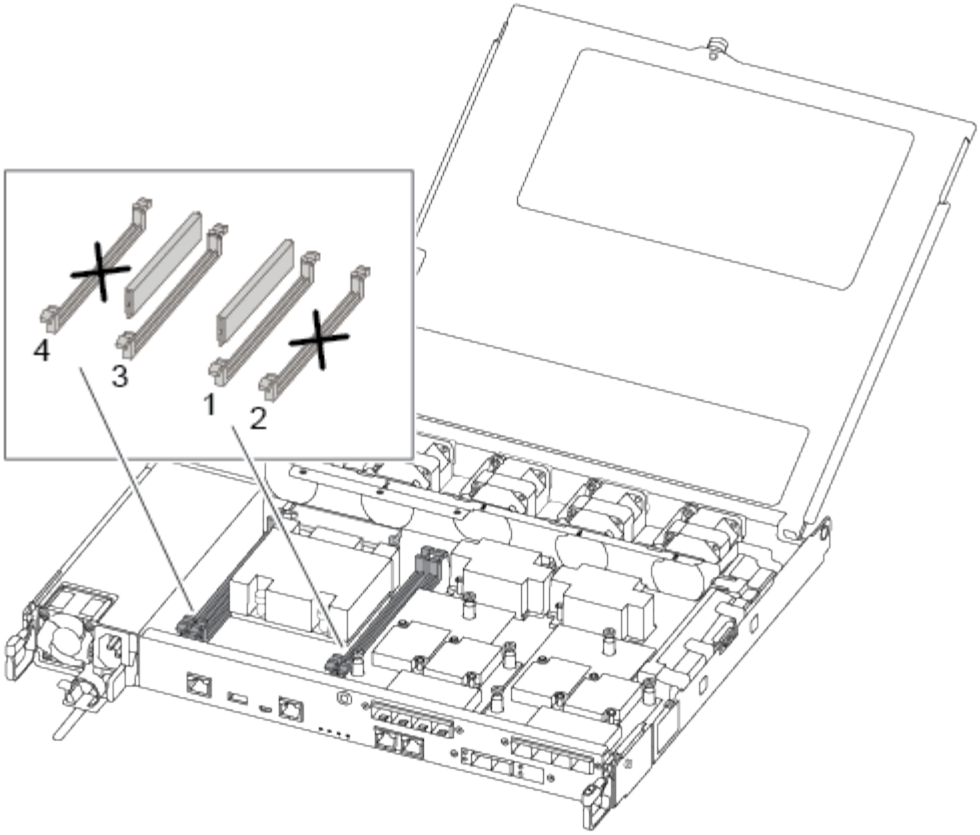

Install each DIMM into the same slot it occupied in the impaired controller module.

1. Slowly push apart the DIMM ejector tabs on either side of the DIMM, and slide the DIMM out of the slot.

   Hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

2. Locate the corresponding DIMM slot on the replacement controller module.
3. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the DIMM squarely into the socket.

   The DIMMs fit tightly in the socket. If not, reinsert the DIMM to realign it with the socket.

4. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.
5. Repeat these steps for the remaining DIMM.

Step 6: Move a mezzanine card

To move a mezzanine card, you must remove the cabling and any QSFPs and SFPs from the ports, move the mezzanine card to the replacement controller, reinstall any QSFPs and SFPs onto the ports, and cable the ports.

1. Locate and move the mezzanine cards from your impaired controller module.

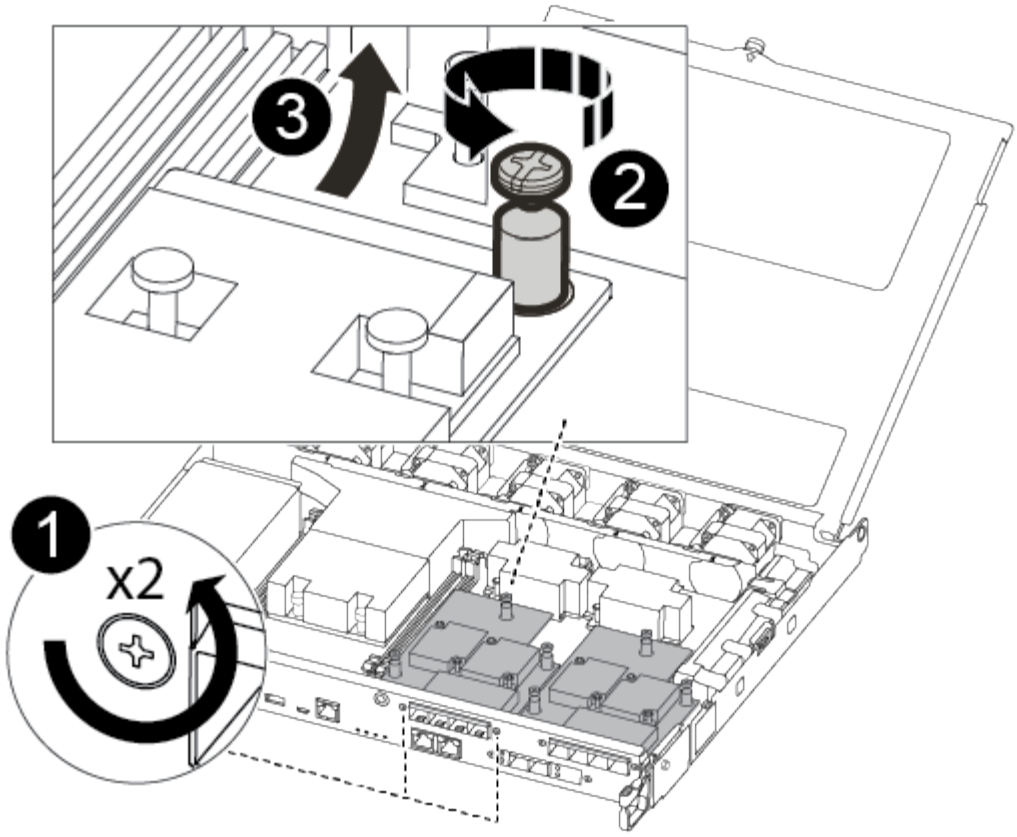

1	Remove screws on the face of the controller module.
2	Loosen the screw in the controller module.
3	Move the mezzanine card.

2. Unplug any cabling associated with the mezzanine card.

   Make sure that you label the cables so that you know where they came from.

   a. Remove any SFP or QSFP modules that might be in the mezzanine card and set it aside.
   b. Using the #1 magnetic screwdriver, remove the screws from the face of the impaired controller module and from the mezzanine card, and set them aside safely on the magnet.
   c. Gently lift the mezzanine card out of the socket and move it to the same position in the replacement controller.
   d. Gently align the mezzanine card into place in the replacement controller.
   e. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the replacement controller module and on the mezzanine card.

      Do not apply force when tightening the screw on the mezzanine card; you might crack it.

3. Repeat these steps if there is another mezzanine card in the impaired controller module.

4. Insert the SFP or QSFP modules that were removed onto the mezzanine card.

Step 7: Move the NV battery

When replacing the controller module, you must move the NV battery from the impaired controller module to the replacement controller module.

1. Locate and move the NVMEM battery from your impaired controller module to the replacement controller module.

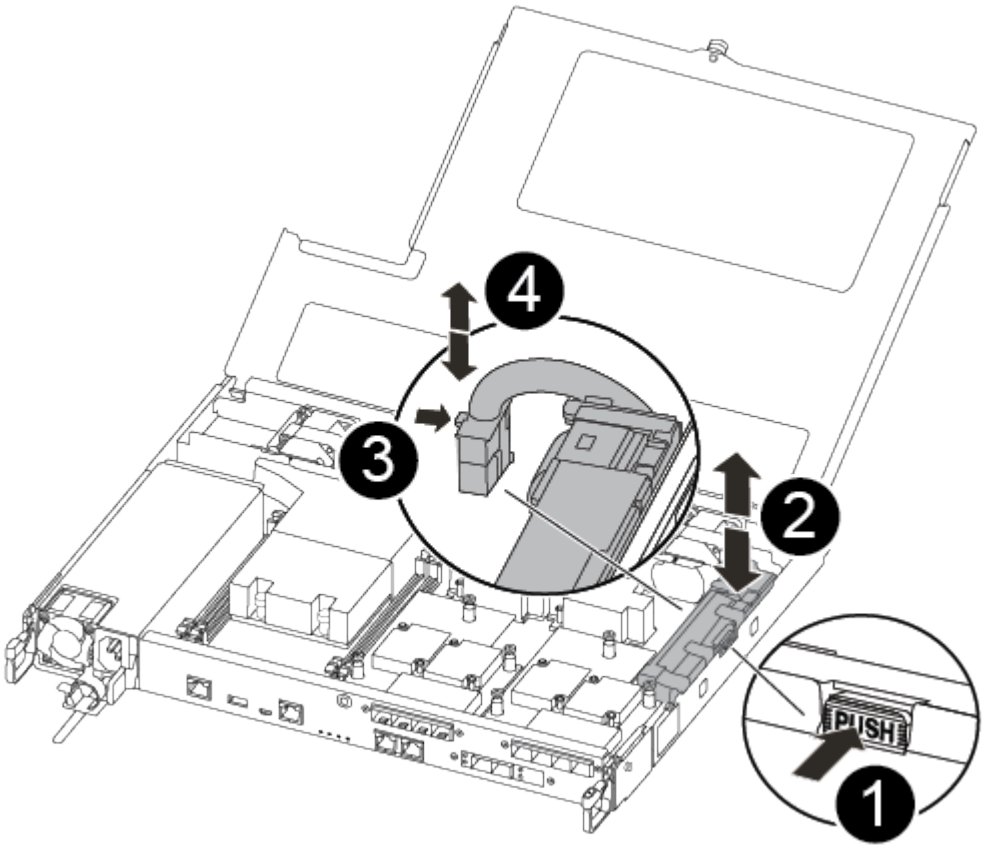

1	Squeeze the clip on the face of the battery plug.
2	Unplug the battery cable from the socket.
3	Grasp the battery and press the blue locking tab marked PUSH.
4	Lift the battery out of the holder and controller module.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.

4. Locate the corresponding NV battery holder on the replacement controller module and align the NV battery to the battery holder.
5. Insert the NV battery plug into the socket.
6. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
7. Press firmly down on the battery pack to make sure that it is locked into place.

Step 8: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

You can use the following illustrations or the written steps to install the replacement controller module in the chassis.

1. If you have not already done so, install the air duct.

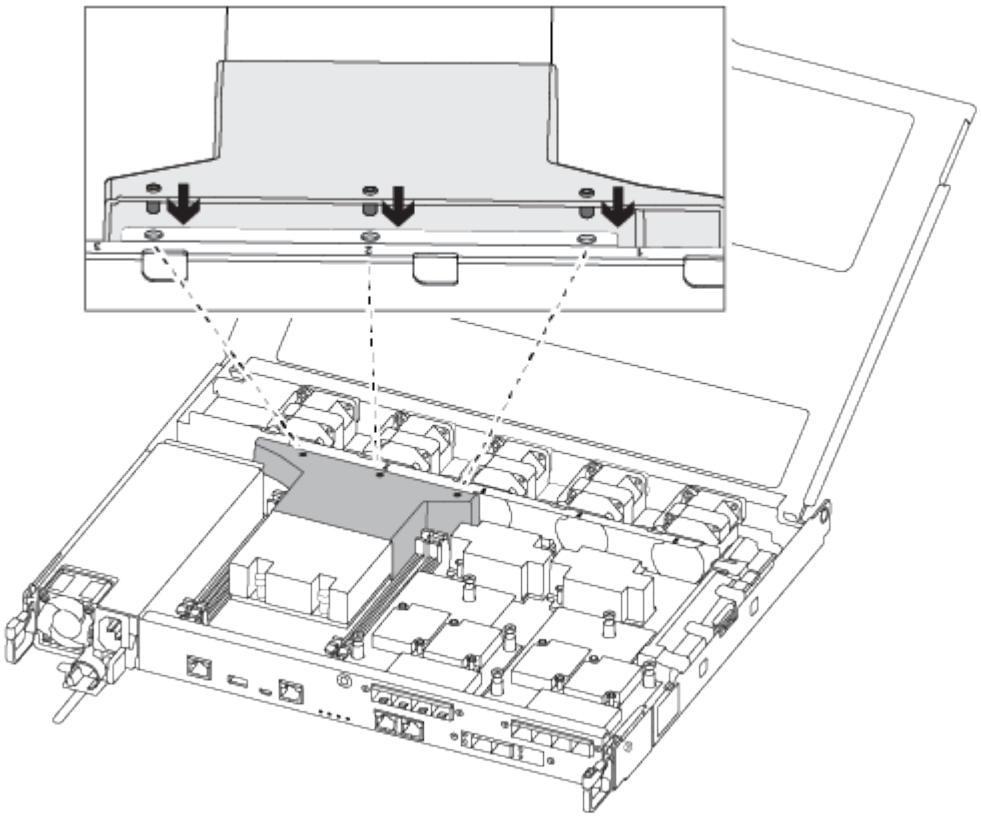

2. Close the controller module cover and tighten the thumbscrew.

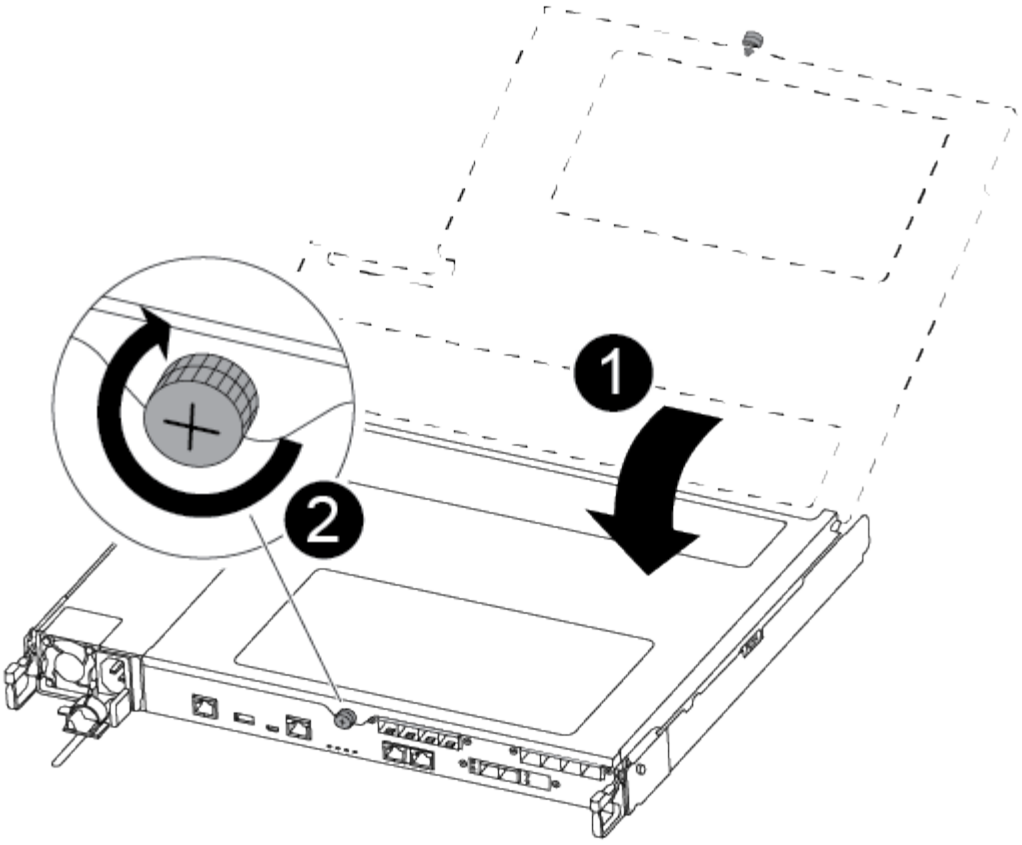

1	Controller module cover
2	Thumbscrew

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching

mechanisms snap into place.

The controller module should be fully inserted and flush with the edges of the chassis.

f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

Restore and verify the system configuration - ASA C250

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - ha
   - mcc
   - mccip
   - non-ha

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - ASA C250

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and

then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - ASA C250

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - ASA C250

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

About this task

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

> If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

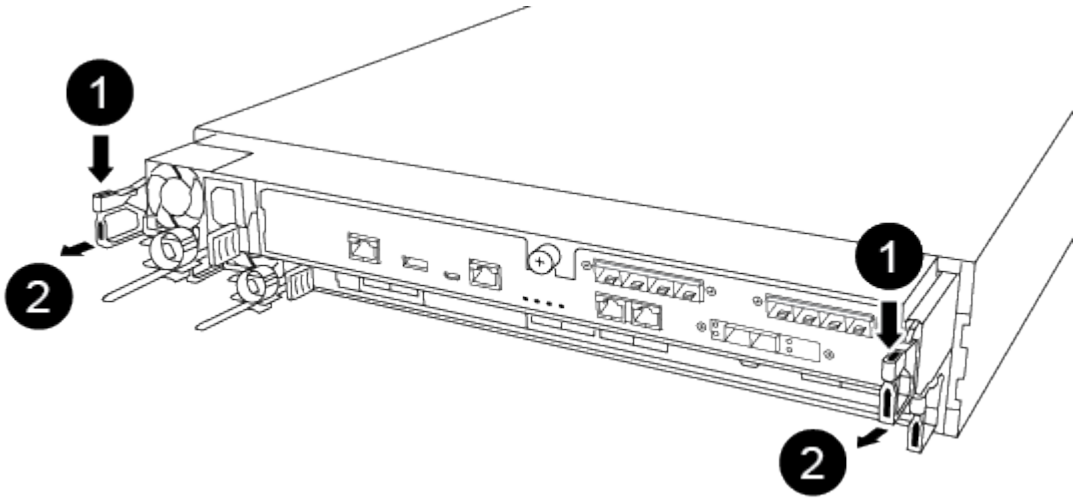

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

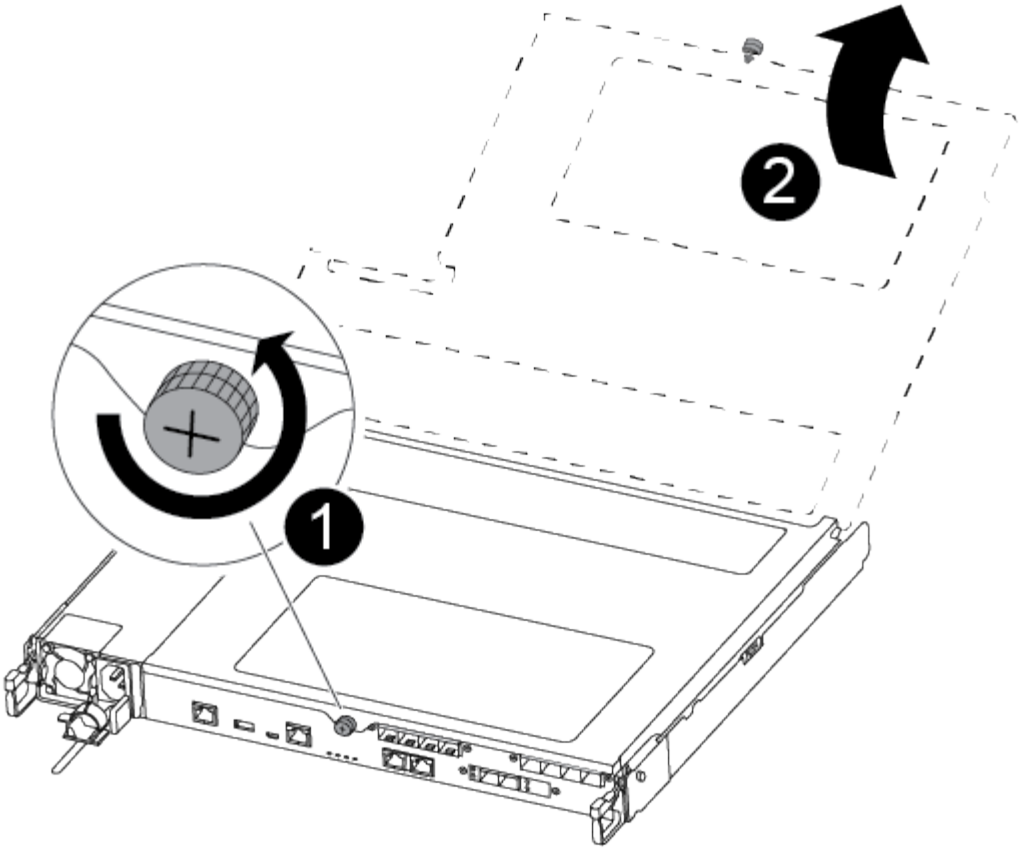

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

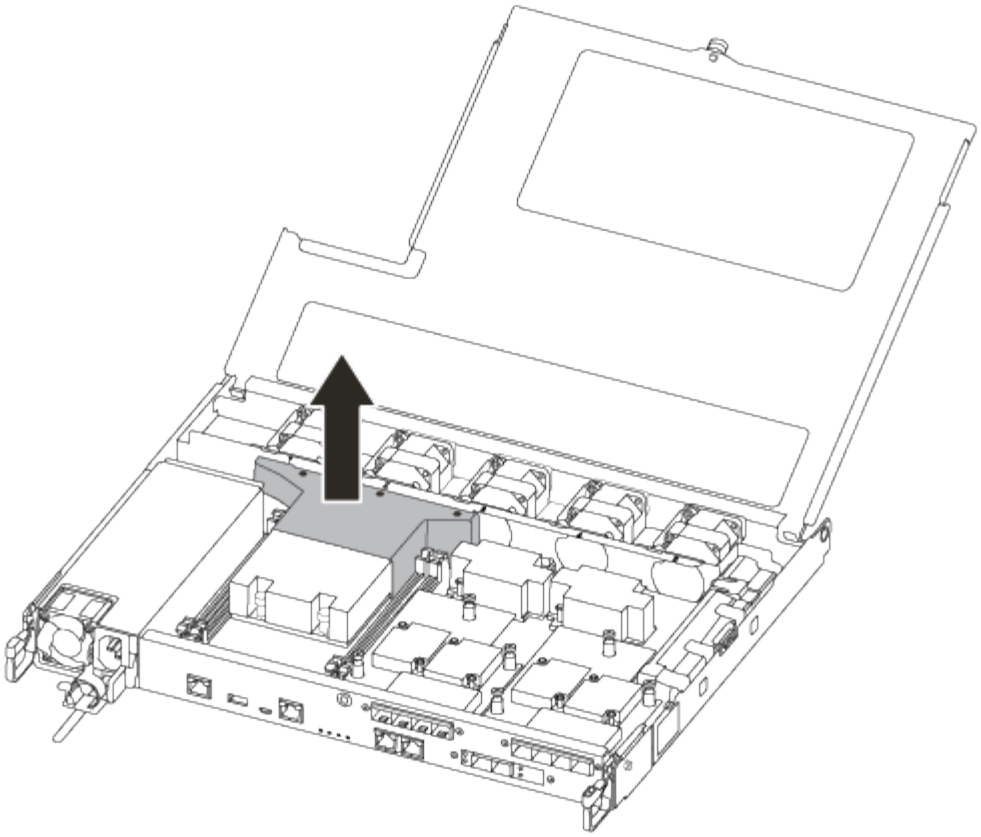

Step 3: Replace a DIMM

To replace a DIMM, you must locate it in the controller module using the DIMM map label on top of the air duct and then replace it following the specific sequence of steps.

Use the following video or the tabulated steps to replace a DIMM:

Animation - Replace a DIMM

1. Replace the impaired DIMM on your controller module.

   The DIMMs are in slot 3 or 1 on the motherboard. Slot 2 and 4 are left empty. Do not attempt to install DIMMs into these slots.

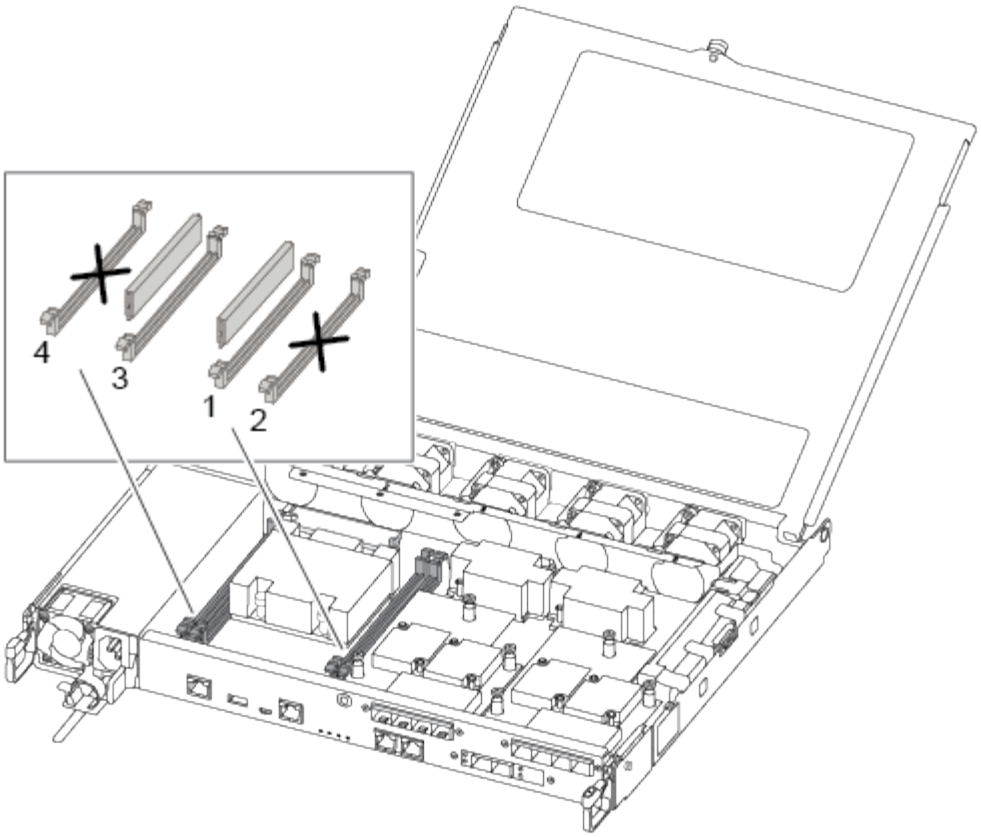

2. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
3. Slowly push apart the DIMM ejector tabs on either side of the DIMM, and slide the DIMM out of the slot.
4. Leave DIMM ejector tabs on the connector in the open position.
5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   Hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

6. Insert the replacement DIMM squarely into the slot.

   The DIMMs fit tightly in the socket. If not, reinsert the DIMM to realign it with the socket.

7. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following illustrations or the written steps to install the replacement controller module in the chassis.

1. If you have not already done so, install the air duct.

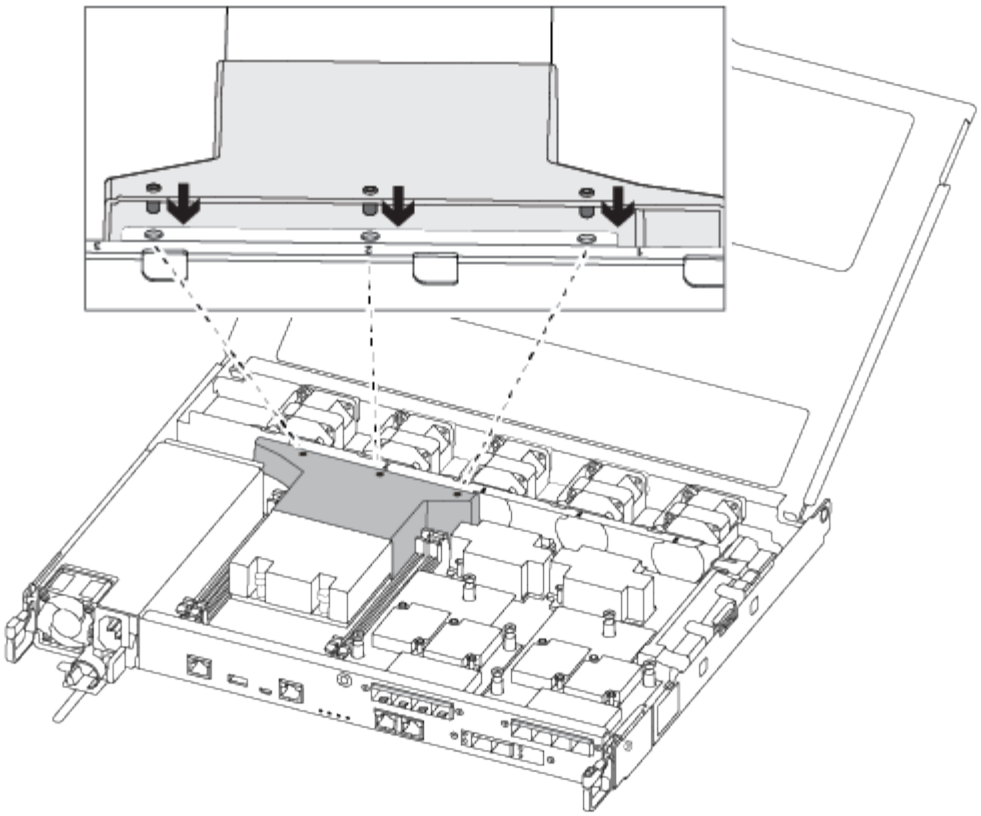

2. Close the controller module cover and tighten the thumbscrew.

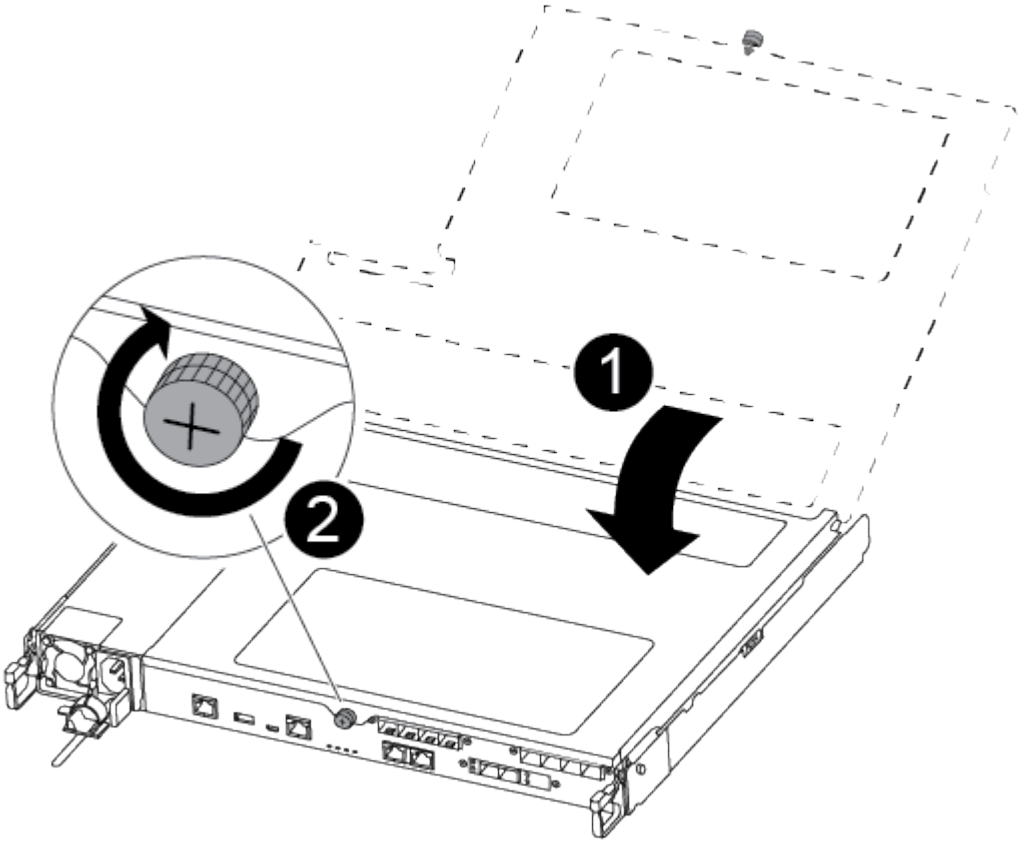

1	Controller module cover
2	Thumbscrew

3. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.
4. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.
5. Recable the system, as needed.
6. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
7. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - ASA C250

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system

console.

The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan - ASA C250

To replace a fan, remove the failed fan module and replace it with a new fan module.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

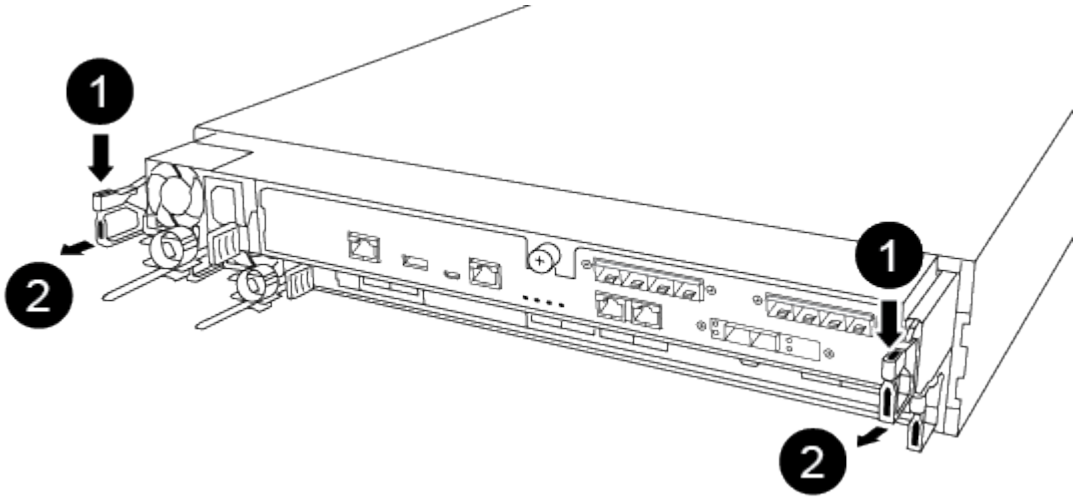

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

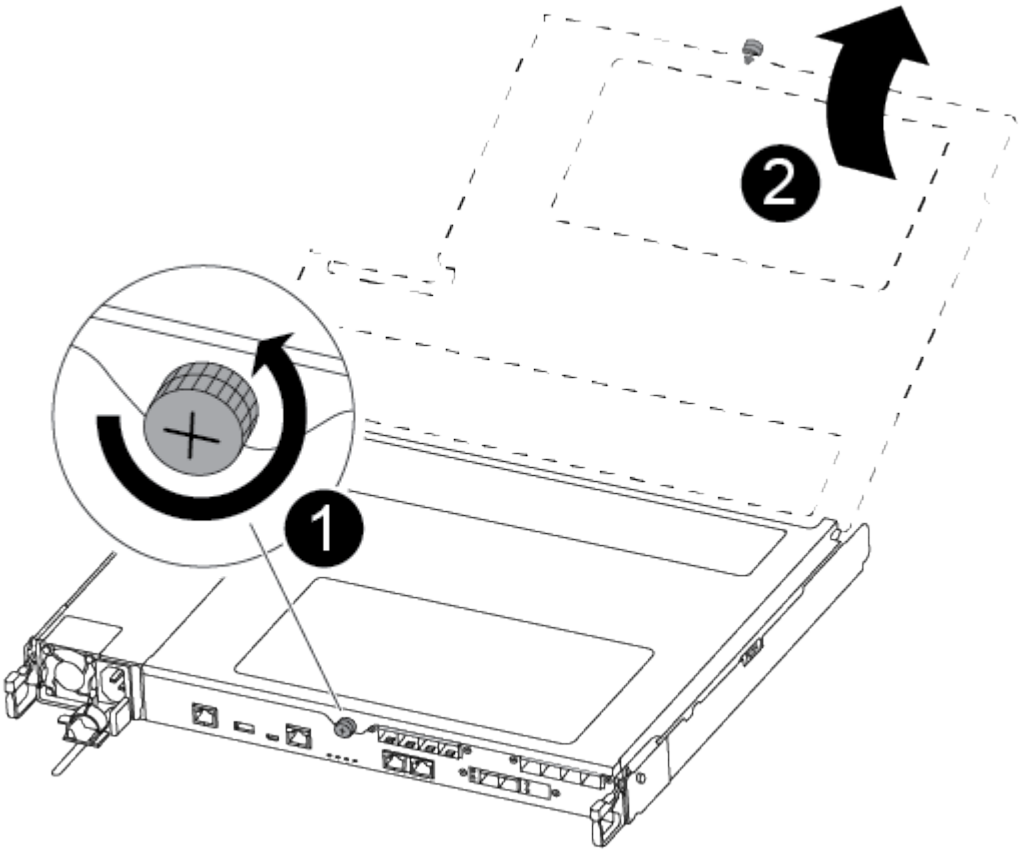

1	Thumbscrew
2	Controller module cover

Step 3: Replace a fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

Use the following video or the tabulated steps to replace a fan:

Animation - Replace a fan

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
2. Remove the fan module by pinching the side of the fan module, and then lifting the fan module straight out of the controller module.

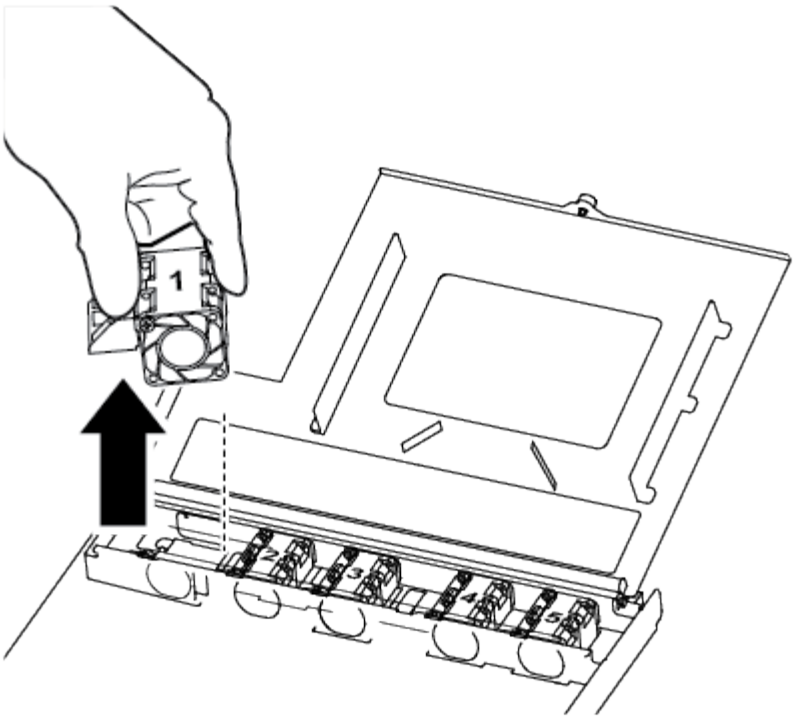

1	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Close the controller module cover and tighten the thumbscrew.

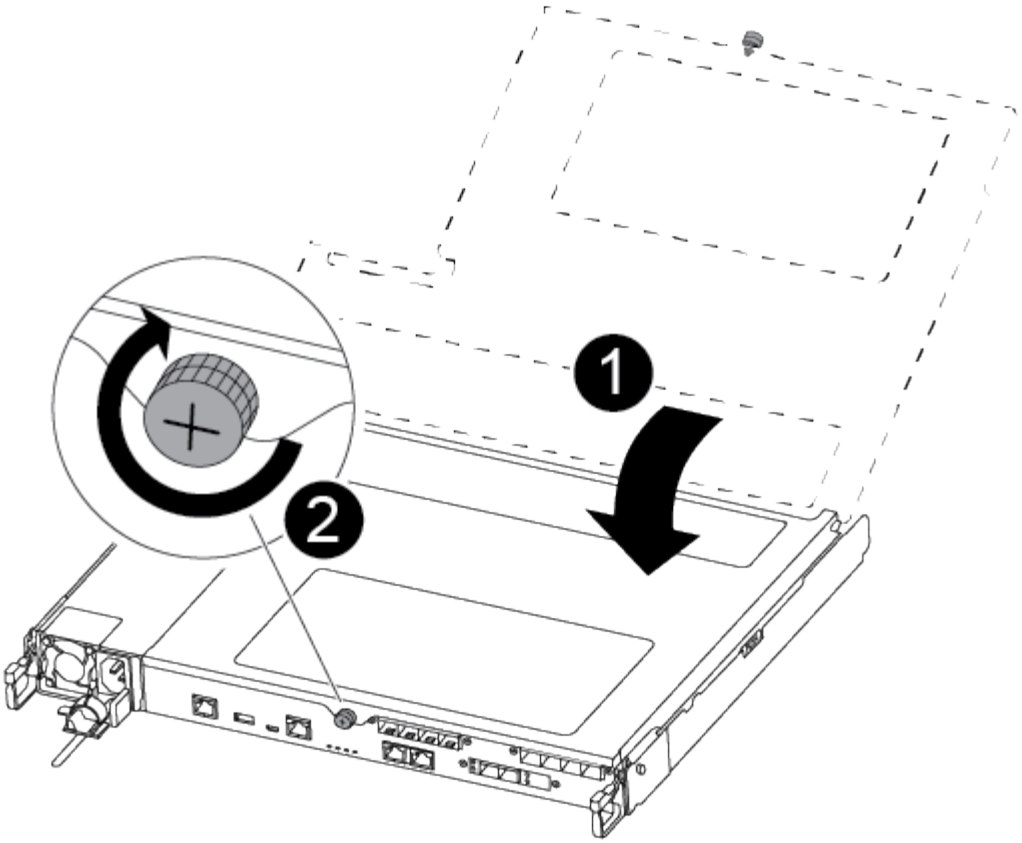

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.

   f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace or install a mezzanine card - ASA C250

To replace a failed mezzanine card, you must remove the cables and any SFP or QSFP modules, replace the card, reinstall the SFP or QSFP modules and recable the cards. To install a new mezzanine card, you must have the appropriate cables and SFP or QSFP modules.

About this task

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

Remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

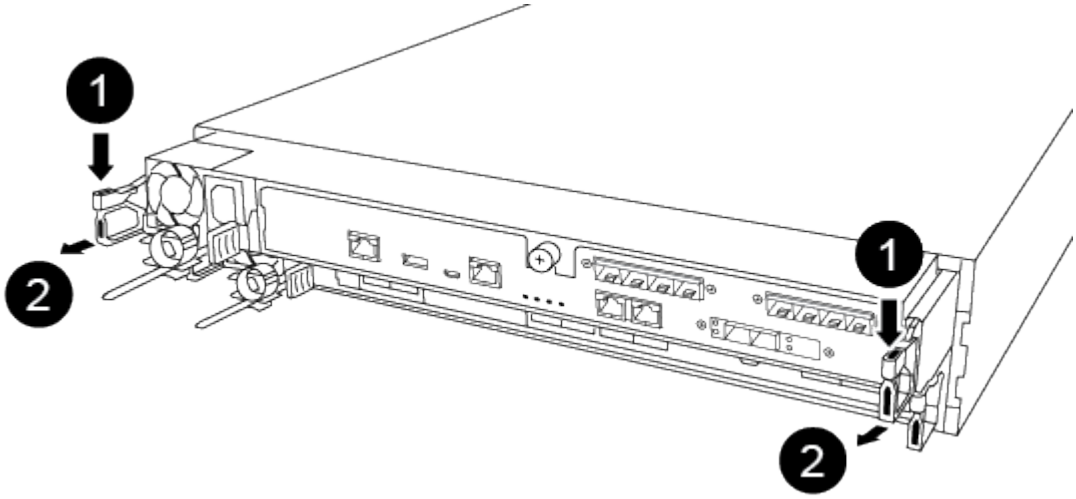

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

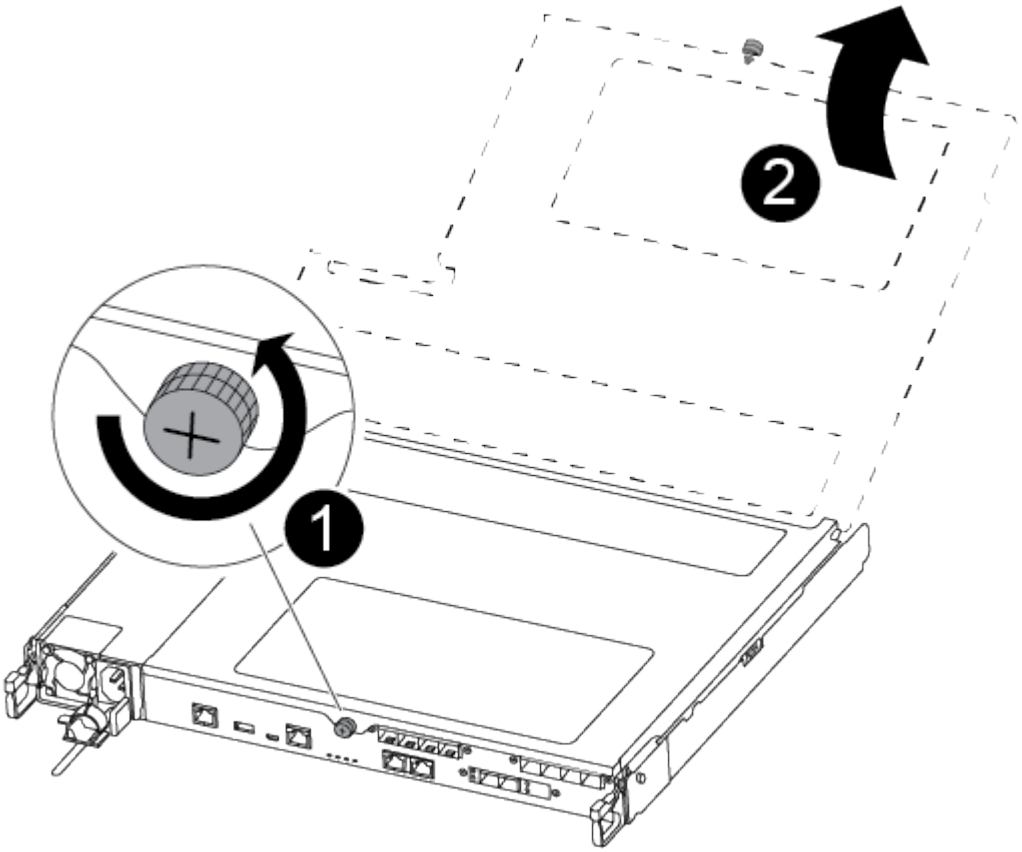

1	Thumbscrew
2	Controller module cover.

Step 3: Replace or install a mezzanine card

To replace a mezzanine card, you must remove the impaired card and install the replacement card; to install a mezzanine card, you must remove the faceplate and install the new card.

Use the following video or the tabulated steps to replace a mezzanine card:

Animation - Replace a mezzanine card

1. To replace a mezzanine card:
2. Locate and replace the impaired mezzanine card on your controller module.

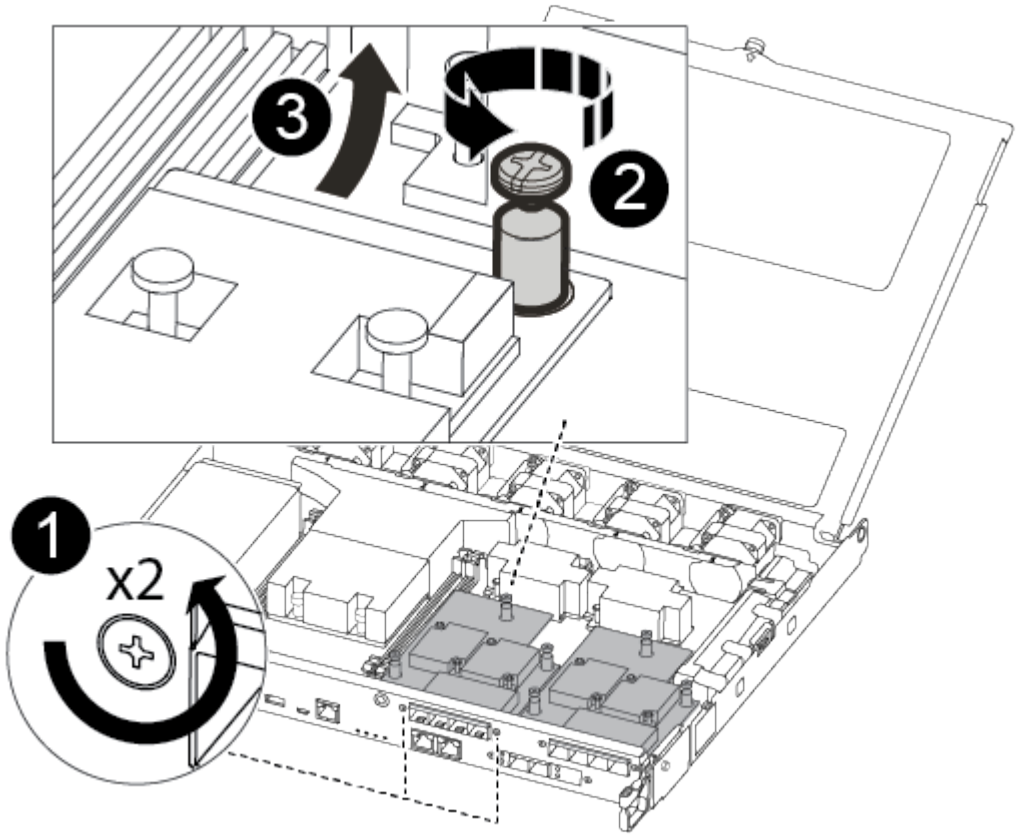

1	Remove screws on the face of the controller module.
2	Loosen the screw in the controller module.
3	Remove the mezzanine card.

a. Unplug any cabling associated with the impaired mezzanine card.

   Make sure that you label the cables so that you know where they came from.

b. Remove any SFP or QSFP modules that might be in the impaired mezzanine card and set it aside.
c. Using the #1 magnetic screwdriver, remove the screws from the face of the controller module and set them aside safely on the magnet.
d. Using the #1 magnetic screwdriver, loosen the screw on the impaired mezzanine card.
e. Using the #1 magnetic screwdriver, gently lift the impaired mezzanine card directly out of the socket and set it aside.
f. Remove the replacement mezzanine card from the antistatic shipping bag and align it to the inside face of the controller module.
g. Gently align the replacement mezzanine card into place.

h. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the controller module and on the mezzanine card.

> Do not apply force when tightening the screw on the mezzanine card; you might crack it.

i. Insert any SFP or QSFP modules that were removed from the impaired mezzanine card to the replacement mezzanine card.

3. To install a mezzanine card:
4. You install a new mezzanine card if your system does not have one.
   a. Using the #1 magnetic screwdriver, remove the screws from the face of the controller module and the faceplate covering the mezzanine card slot, and set them aside safely on the magnet.
   b. Remove the mezzanine card from the antistatic shipping bag and align it to the inside face of the controller module.
   c. Gently align the mezzanine card into place.
   d. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the controller module and on the mezzanine card.

> Do not apply force when tightening the screw on the mezzanine card; you might crack it.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Close the controller module cover and tighten the thumbscrew.

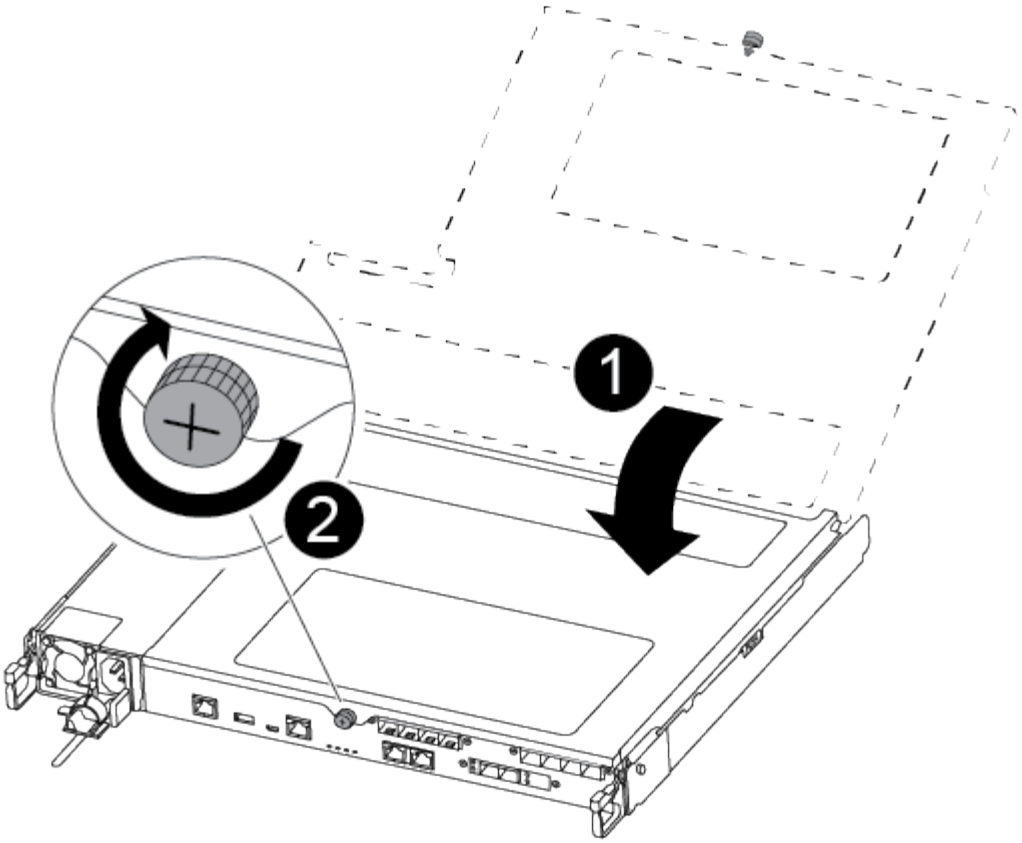

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module should be fully inserted and flush with the edges of the chassis.

   f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVMEM battery - ASA C250

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

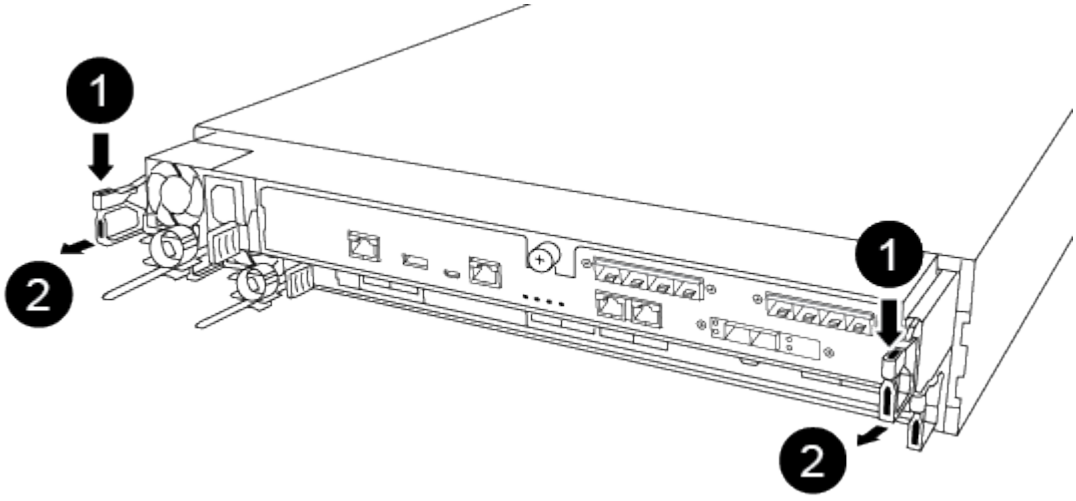

1	Lever

2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

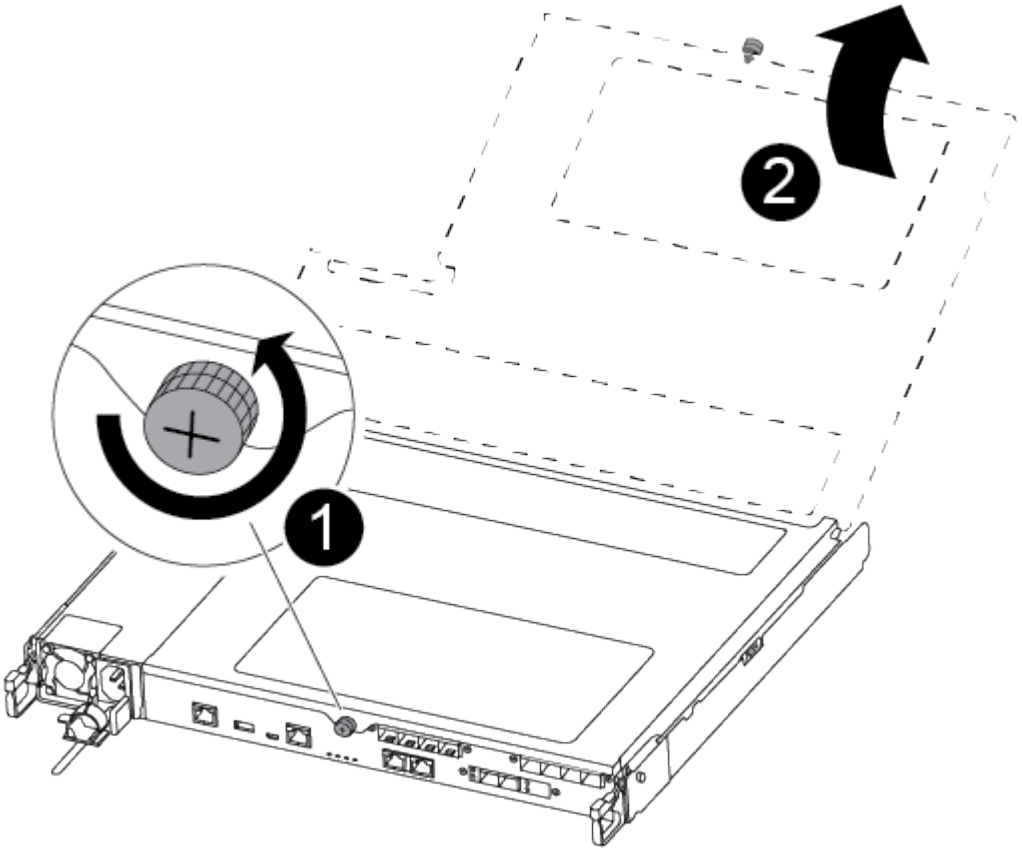

1	Thumbscrew
2	Controller module cover.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module.

Use the following video or the tabulated steps to replace the NVMEM battery:

Animation - Replace the NVMEM battery

1. Locate and replace the impaired NVMEM battery on your controller module.

It is recommended that you follow the illustrated instructions in the order listed.

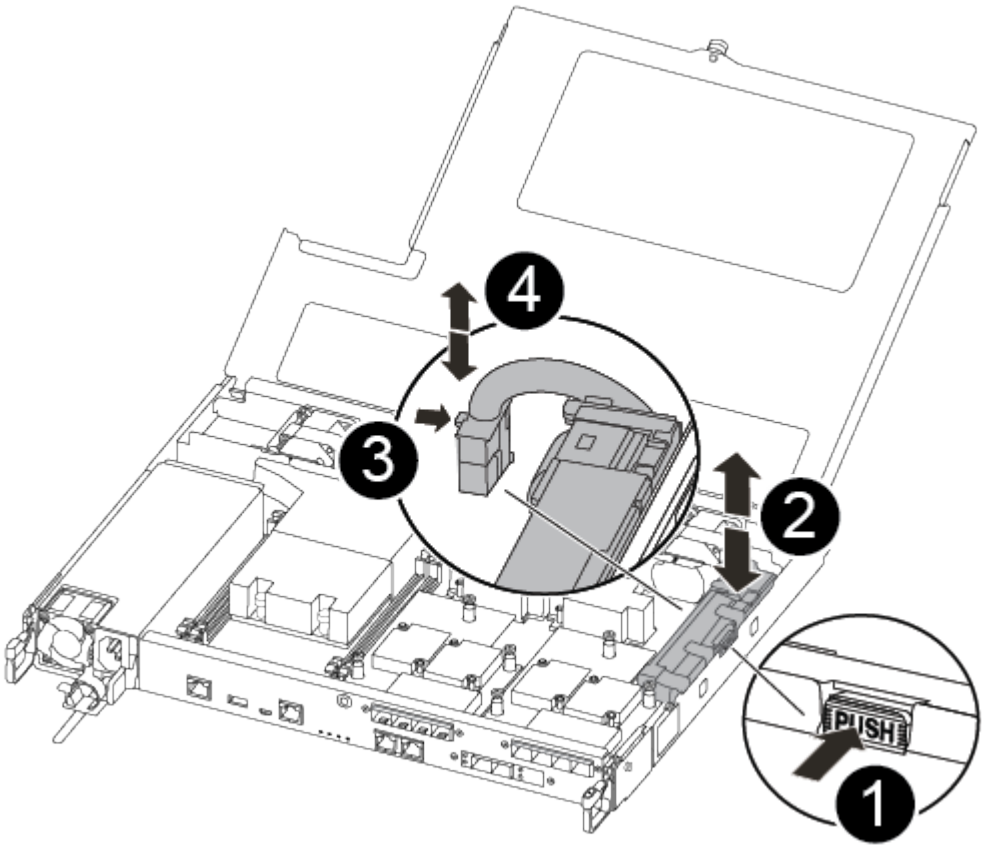

1	Squeeze the clip on the face of the battery plug.
2	Unplug the battery cable from the socket.
3	Grasp the battery and press the blue locking tab marked PUSH.
4	Lift the battery out of the holder and controller module.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module and set it aside.
4. Remove the replacement NV battery from the antistatic shipping bag and align it to the battery holder.
5. Insert the replacement NV battery plug into the socket.
6. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
7. Press firmly down on the battery pack to make sure that it is locked into place.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following illustration or the written steps to install the replacement controller module in the chassis.

1. Close the controller module cover and tighten the thumbscrew.

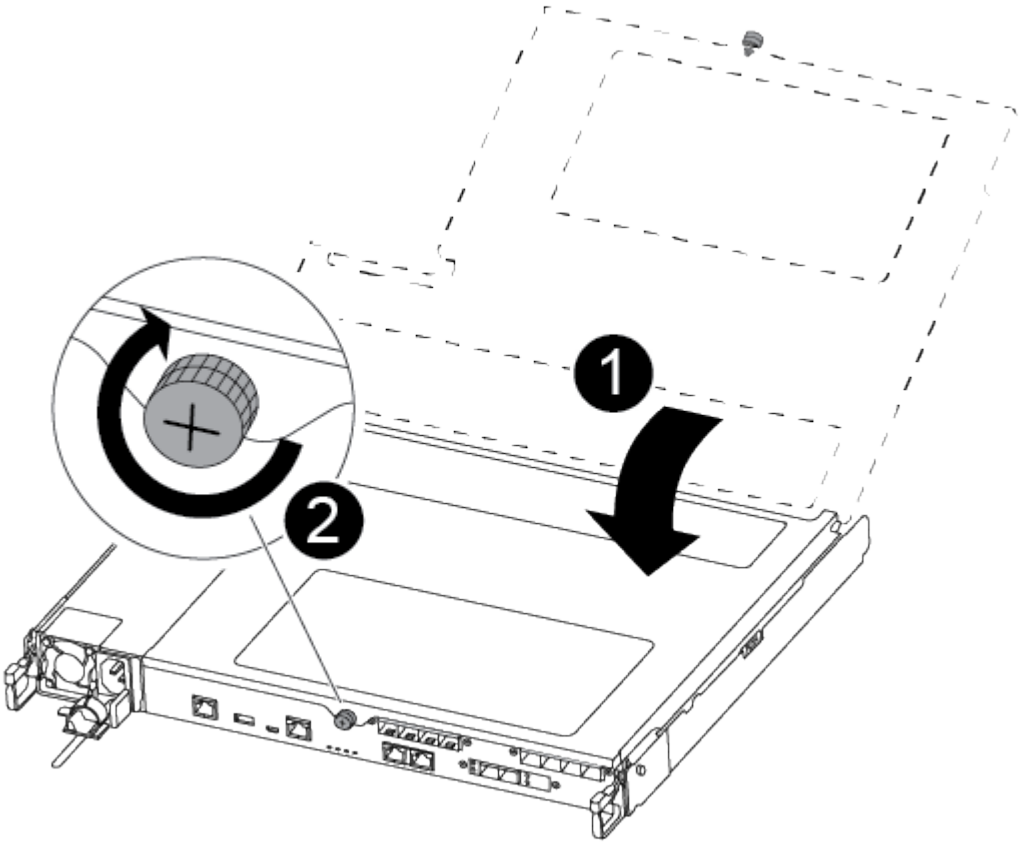

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

The controller module should be fully inserted and flush with the edges of the chassis.

f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

3. Recable the system, as needed.
4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - ASA C250

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one PSU at a time.

> It is a best practice to replace the PSU within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded PSU until the PSU is replaced.

- Power supplies are auto-ranging.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

Use the appropriate procedure for your type of PSU; AC or DC.

Option 1: Replace an AC PSU

Use the following video or the tabulated steps to replace the PSU:

Animation - Replace the AC PSU

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

> The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

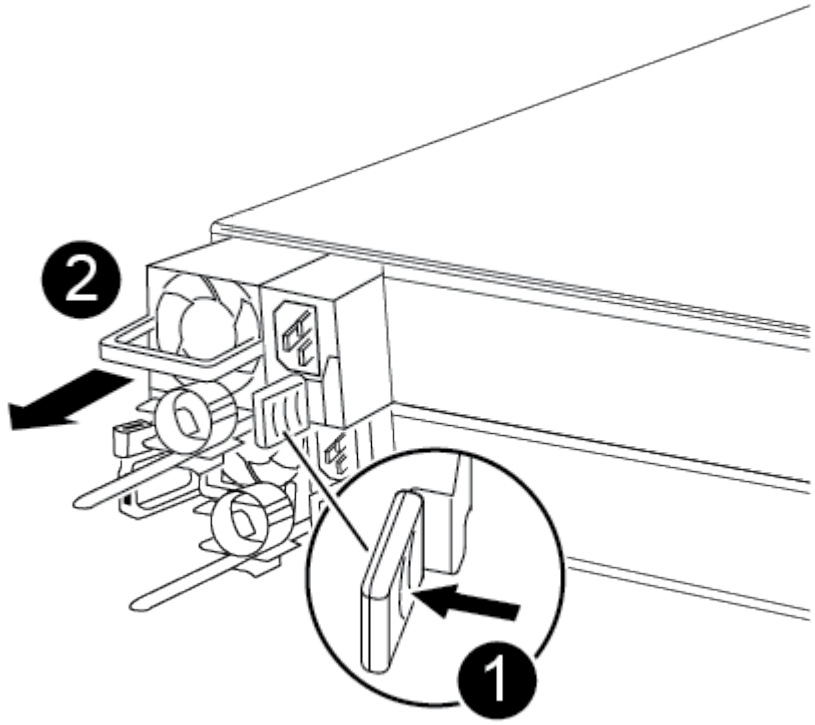

1	Blue PSU locking tab
2	Power supply

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller module.

b. Gently push the PSU into the controller module until the locking tab clicks into place.

The power supplies will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:

a. Reconnect the power cable to the PSU.

b. Secure the power cable to the PSU using the power cable retainer.

Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC power cable connector using the thumb screws on the plug.
   b. Unplug the power cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

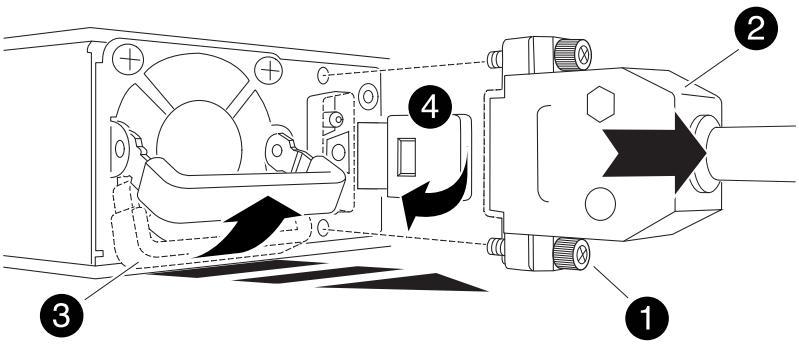

1	Thumb screws
2	D-SUB DC power cable connector

3	Power supply handle
4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA C250

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

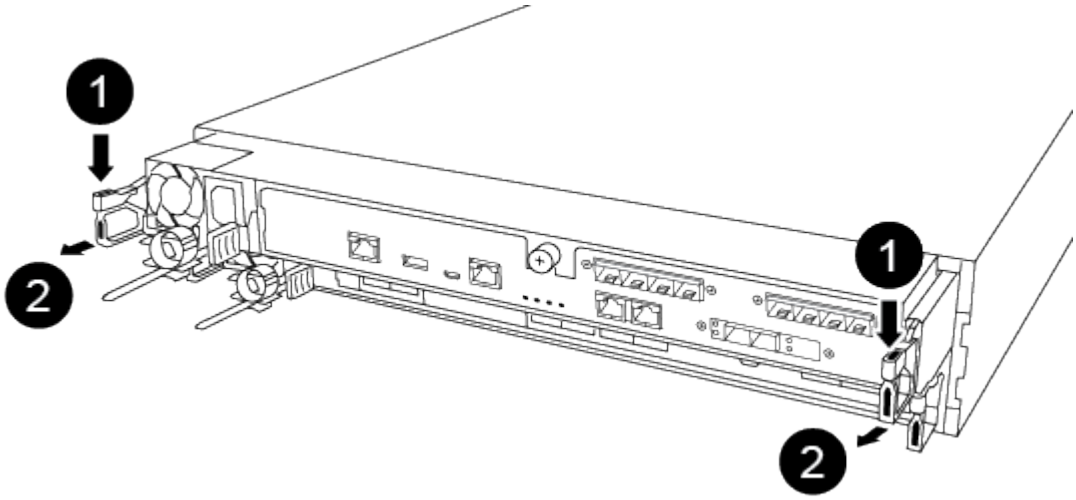

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

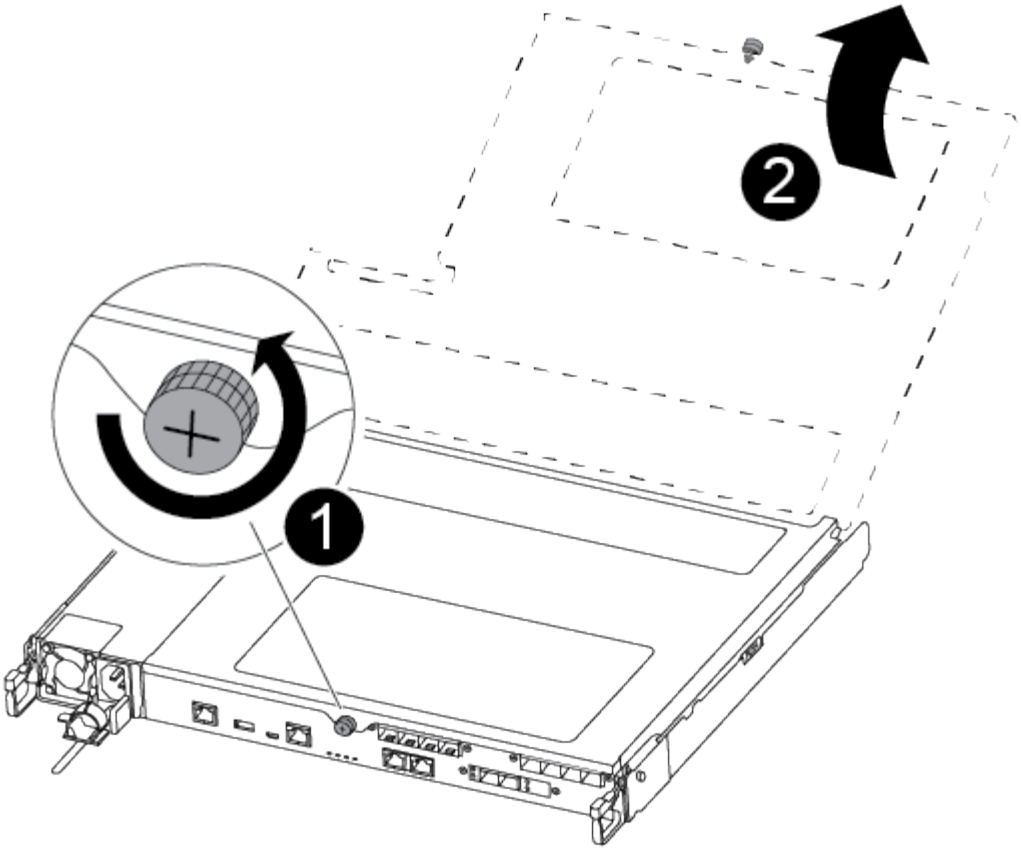

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

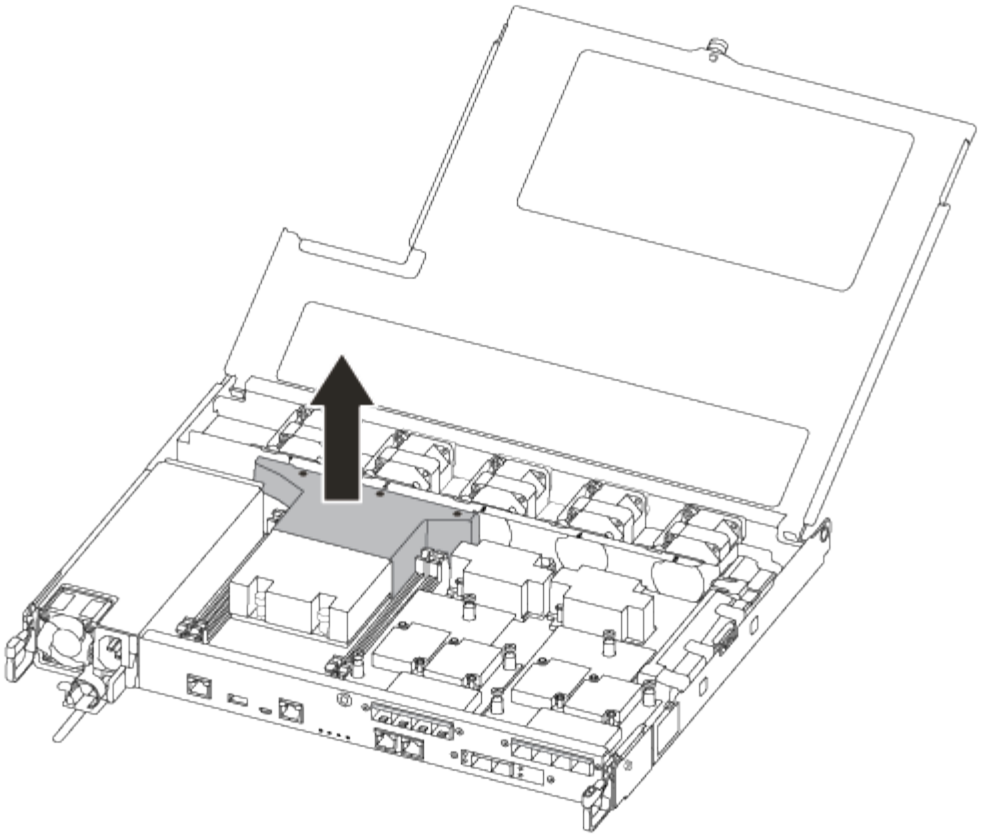

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

Use the following video or the tabulated steps to replace the RTC battery:

Animation - Replace the RTC battery

1. Locate the RTC battery between the heatsink and the midplane and remove it exactly as shown in the graphic.

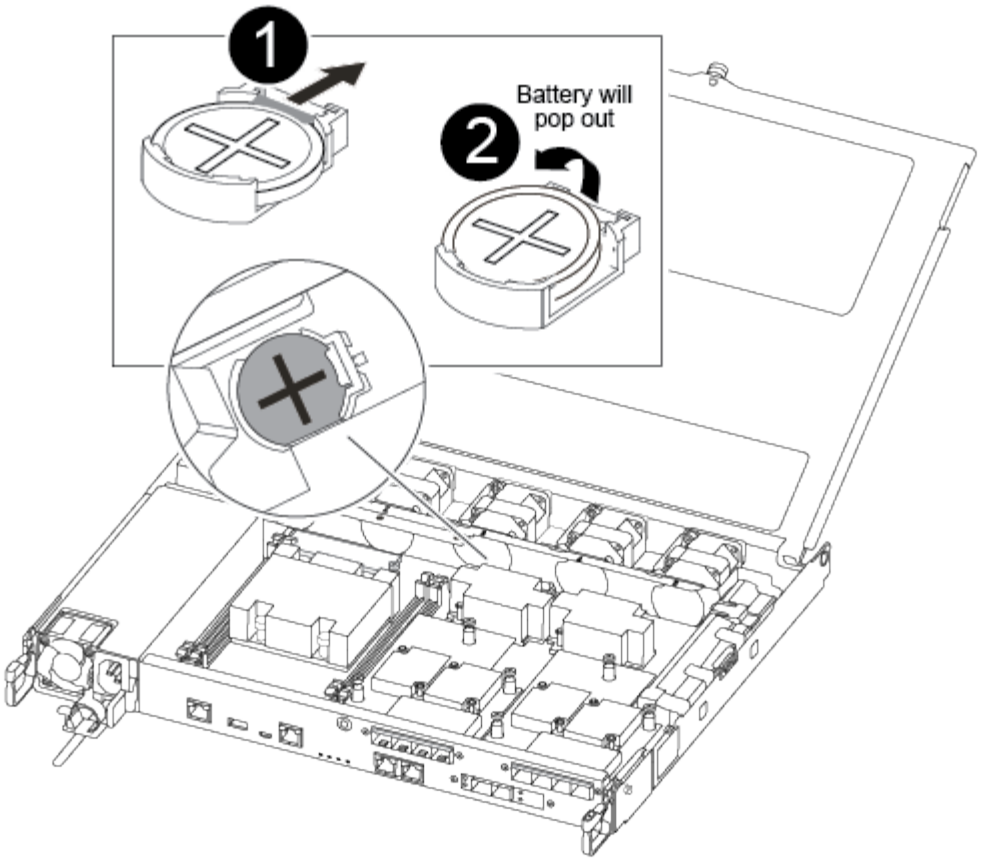

1	Gently pull tab away from the battery housing. **Attention:** Pulling it away aggressively might displace the tab.
2	Lift the battery up. **Note:** Make a note of the polarity of the battery.
3	The battery should eject out.

2. Remove the replacement battery from the antistatic shipping bag.
3. Locate the RTC battery holder between the heatsink and the midplane and insert it exactly as shown in the graphic.

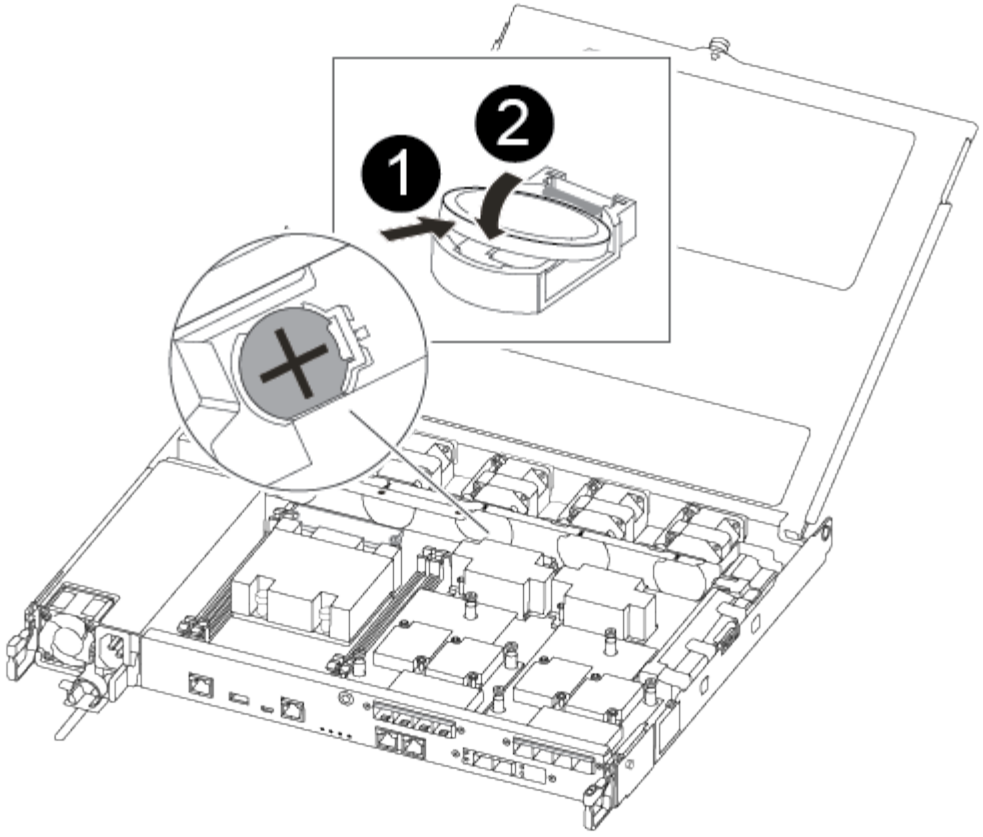

1	With positive polarity face up, slide the battery under the tab of the battery housing.
2	Push the battery gently into place and make sure the tab secures it to the housing. Pushing it in aggressively might cause the battery to eject out again.

4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Insert the controller module into the chassis:
    a. Ensure the latching mechanism arms are locked in the fully extended position.
    b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
    c. Place your index fingers through the finger holes from the inside of the latching mechanism.
    d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
    e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

       The controller module should be fully inserted and flush with the edges of the chassis.
    f. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

       The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.
    g. Halt the controller at the LOADER prompt.
5. Reset the time and date on the controller:
    a. Check the date and time on the healthy controller with the `show date` command.
    b. At the LOADER prompt on the target controller, check the time and date.
    c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
    d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
    e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA C400 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

For MetroCluster configurations, see either:

- Install MetroCluster IP configuration
- Install MetroCluster Fabric-Attached configuration

Quick guide - ASA C400

The quick guide provides graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this procedure if you are familiar with installing NetApp systems.

Use the AFF C400 Installation and Setup Instructions.

> The ASA C400 uses the same installation procedure as the AFF C400 system.

Videos - ASA C400

The following video shows how to install and cable your new system.

Animation - AFF C400 Installation and setup instructions

> The ASA C400 uses the same installation procedure as the AFF C400 system.

If you have a MetroCluster configuration, use the MetroCluster installation content.

MetroCluster Documentation

Detailed guide - ASA C400

This guide gives detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

If you have a MetroCluster configuration, use the MetroCluster installation content.

MetroCluster Documentation

Step 1: Prepare for installation

To install your system, you need to create an account, register the system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

- You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

  NetApp Hardware Universe

  Find the Release Notes for your version of ONTAP 9

- You need to provide the following at your site:
  - Rack space for the storage system
  - Phillips #2 screwdriver
  - Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
100 GbE cable (QSFP28)	X66211A-05 (112-00595), 0.5m X66211A-1 (112-00573), 1m X66211A-2 (112-00574), 2m X66211A-5 (112-00574), 5m		Storage, cluster interconnect/HA, and Ethernet data (order-dependent)
25 GbE cable (SFP28)	X66240-2 (112-00598), 2m X66240-5 (112-00639), 5m		GbE network connection (order-dependent)
32 Gb FC (SFP+ Op)	X66250-2 (112-00342), 2m X66250-5 (112-00344), 5m X66250-15 (112-00346), 15m		FC network connection

Type of cable…	Part number and length	Connector type	For…
Optical cables	X66250-2-N-C (112-00342)		16 Gb FC or 25GbE cables for mezzanine cards (order-dependent)
RJ-45 (order dependent)	X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network
Micro-USB console cable	Not applicable		Console connection used during software setup if laptop or console does not support network discovery.
Power cables	Not applicable		Powering up the system

4. Review the *NetApp ONTAP Configuration Guide* and collect the required information listed in that guide.

   ONTAP Configuration Guide

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

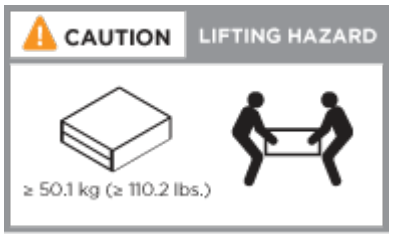

3. Attach cable management devices (as shown).

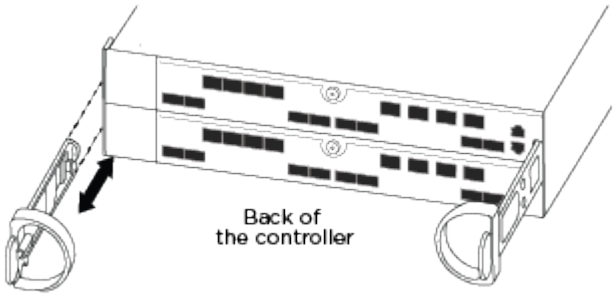

4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the switched cluster method.

About this task

- If the port labels on the card are not visible, you can identify the ports by checking the card installation orientation (for C400, the PCIe connector socket is on the left side of the card slot), and then look for the card by part number in NetApp Hardware Universe, which shows a graphic of the bezel with the port labels. You can find the card part number using the sysconfig -a command or on the system packing list.
- If you are cabling an MetroCluster IP configuration, ports e0a/e0b are available for hosting data LIFs (usually in Default IPSpace).

Option 1: Cable a two-node switchless cluster

A controller module's cluster interconnect and HA ports are cabled to its partner controller module. The optional data ports, optional NIC cards, and management ports on the controller modules are connected to switches.

Before you begin

You must have contacted your network administrator for information about connecting the system to the switches.

About this task

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.

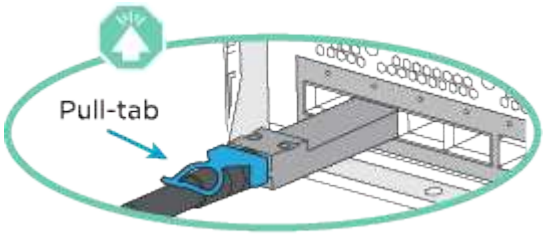

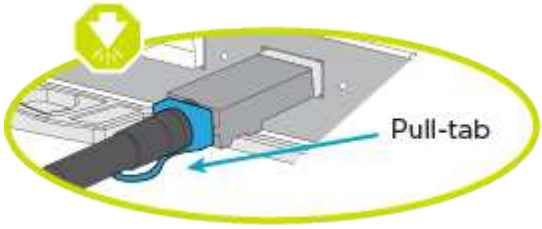

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the illustration to complete the cabling between the controllers and the switches:

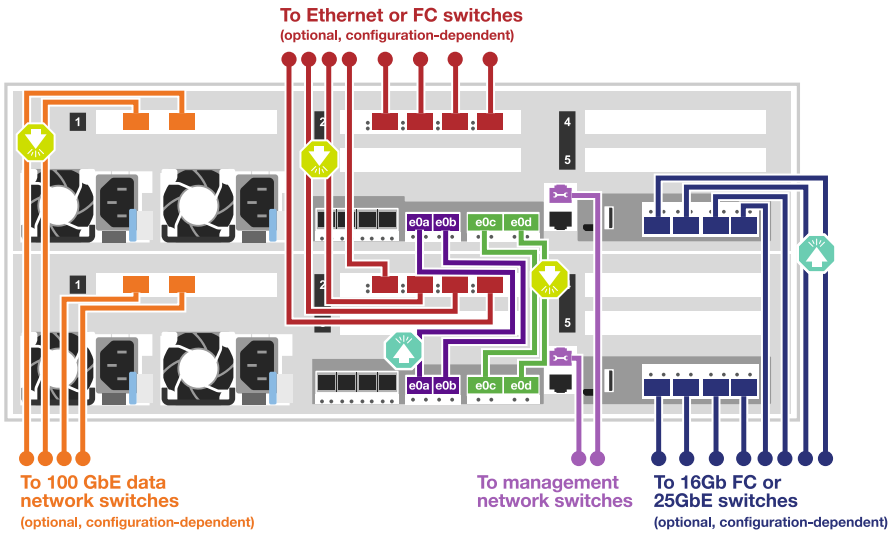

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Option 2: Cable a switched cluster

Controller module cluster interconnect and HA ports are cabled to the cluster/HA switch. The optional data ports, optional NIC cards, mezzanine cards, and management ports are connected to switches.

Before you begin

You must have contacted your network administrator for information about connecting the system to the switches.

About this task

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.

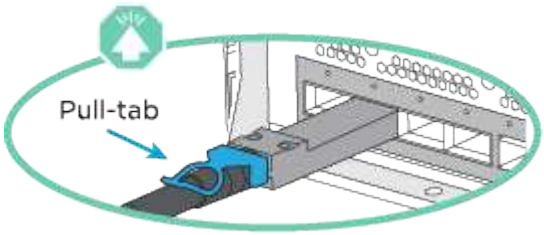

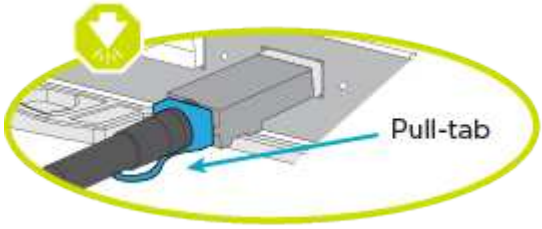

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the illustration to complete the cabling between the controllers and the switches:

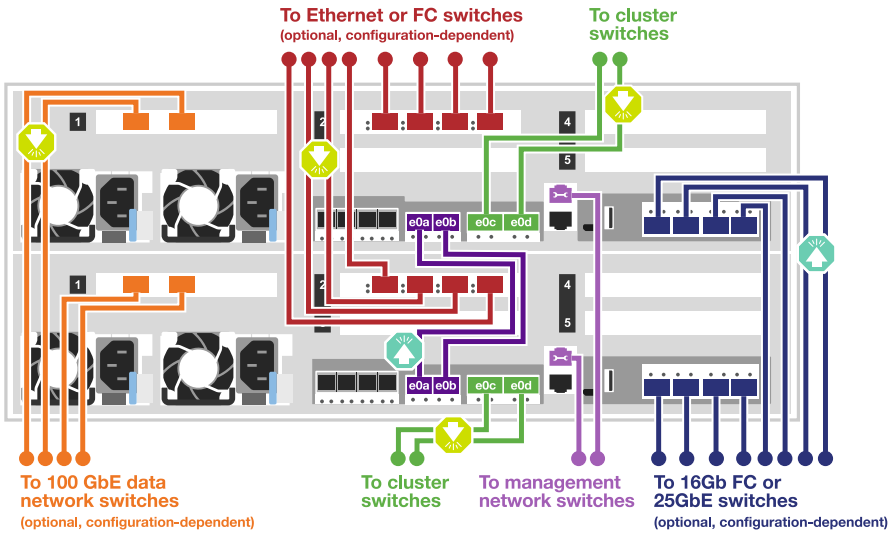

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Step 4: Cable controllers to drive shelves

The following options show you how to cable one or two NS224 drive shelves to your system.

Option 1: Cable the controllers to a single drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

About this task

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the NS224 are up.

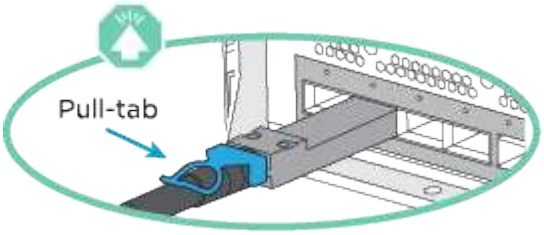

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following illustration to cable your controllers to a single drive shelf.

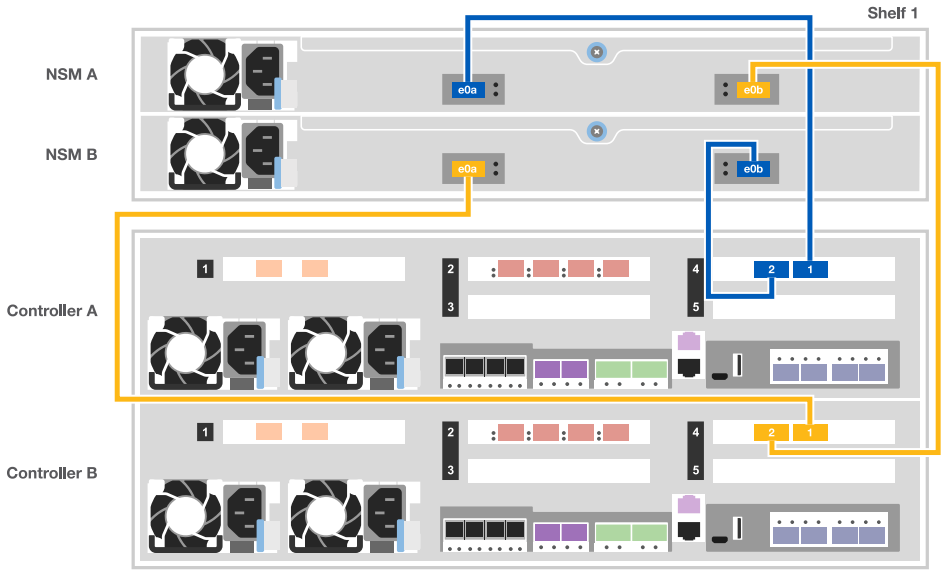

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Option 2: Cable the controllers to two drive shelves

You must cable each controller to the NSM modules on both NS224 drive shelves.

About this task

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the NS224 are up.

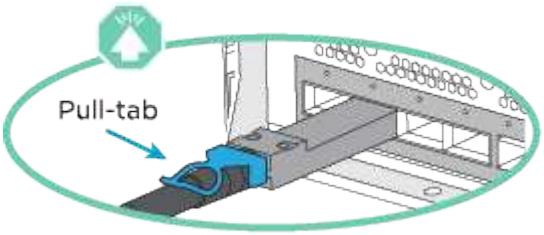

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following illustration to cable your controllers to two drive shelves.

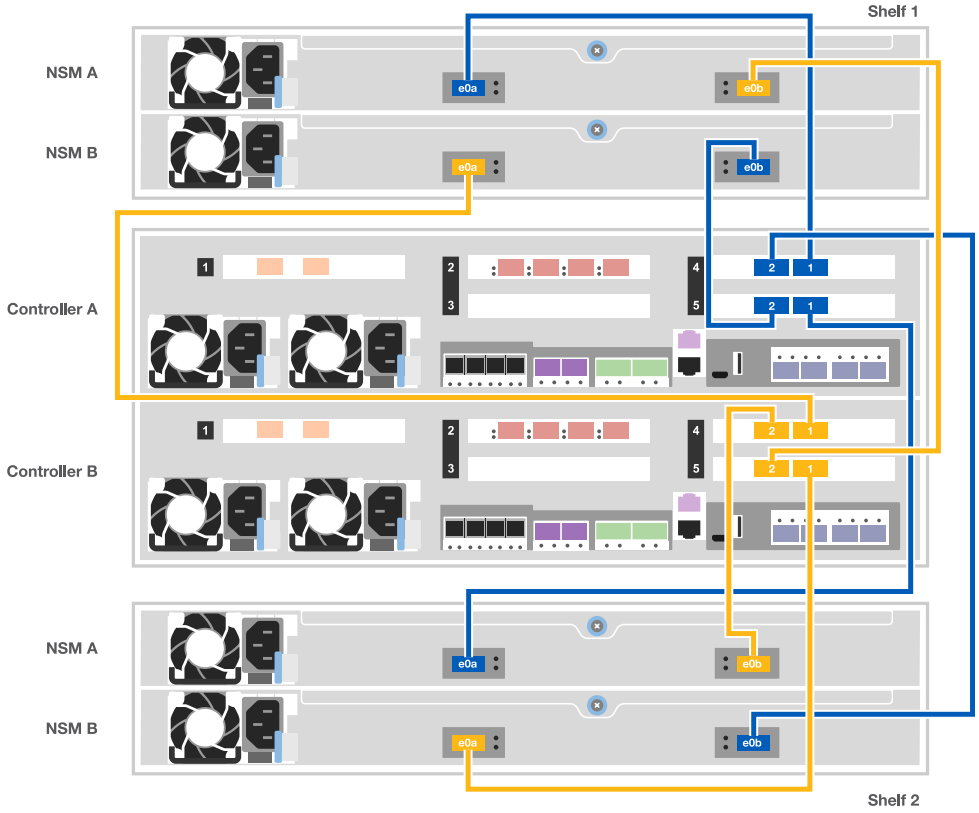

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Completing system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

1. Use the following animation to power on and set shelf IDs for one or more drive shelves:

   For NS224 drive shelves, shelf IDs are pre-set to 00 and 01. If you want to change the shelf IDs, use the straightened end of a paperclip, or narrow tipped ball point pen to access the shelf ID button behind the faceplate.

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Connect your laptop to the Management switch.

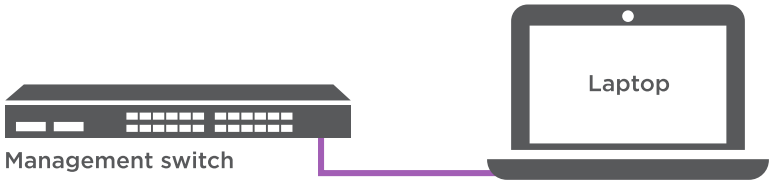

5. Select an ONTAP icon listed to discover:

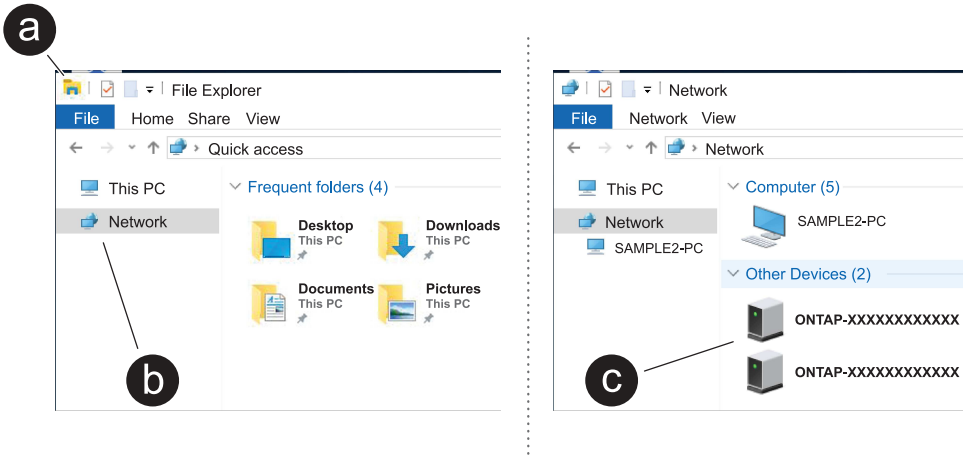

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

   XXXXX is the system serial number for the target node.

   System Manager opens.

6. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

   ONTAP Configuration Guide

7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet .
   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Use the following animation to power on and set shelf IDs for one or more drive shelves:

   For NS224 drive shelves, shelf IDs are pre-set to 00 and 01. If you want to change the shelf IDs, use the straightened end of a paperclip, or narrow tipped ball point pen to access the shelf ID button behind the faceplate.

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   Initial booting may take up to eight minutes.

4. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.

If the management network has DHCP…	Then…
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the *NetApp ONTAP Configuration guide*.

      ONTAP Configuration Guide

6. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain ASA C400 hardware

Maintain the hardware of your ASA C400 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the ASA C400 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the ASA C400 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.
NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.
PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - ASA C400

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your ASA C400 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1 **Review the boot media requirements**

Review the requirements for boot media replacement.

2 **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3 **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4 **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5 **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - ASA C400

Before replacing the boot media in your ASA C400, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.

- /cfcard/kmip/certs/client.crt file.
- /cfcard/kmip/certs/client.key file.
- /cfcard/kmip/certs/CA.pem file.

- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - ASA C400

Shut down the impaired controller in your ASA C400 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - ASA C400

The boot media in your ASA C400 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

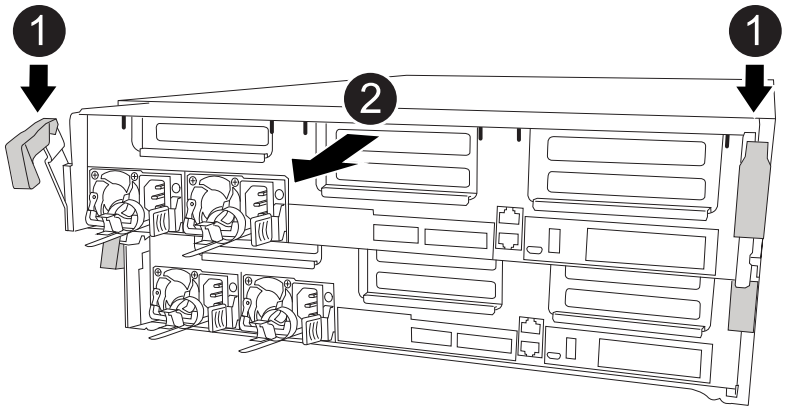

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. Open the air duct:

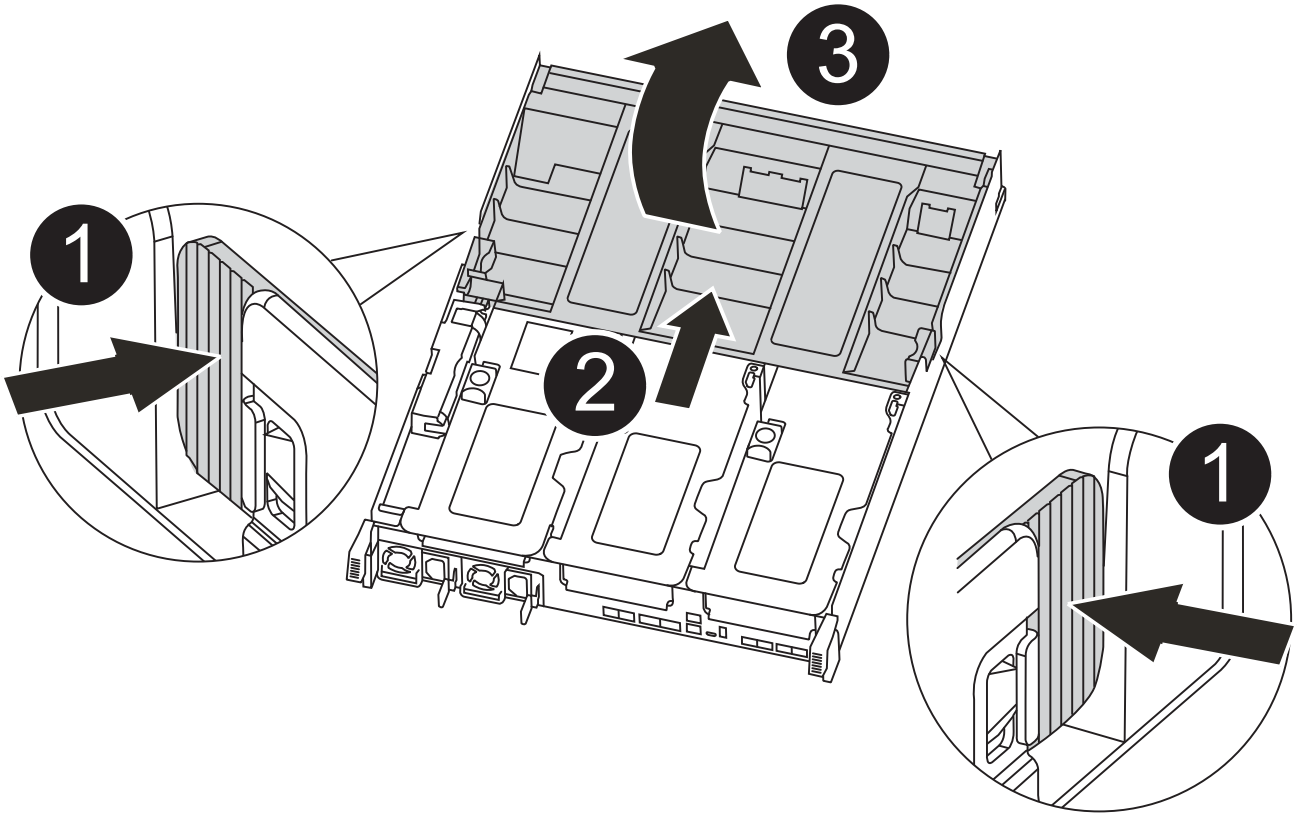

1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.

b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

9. Locate and remove the boot media from the controller module:

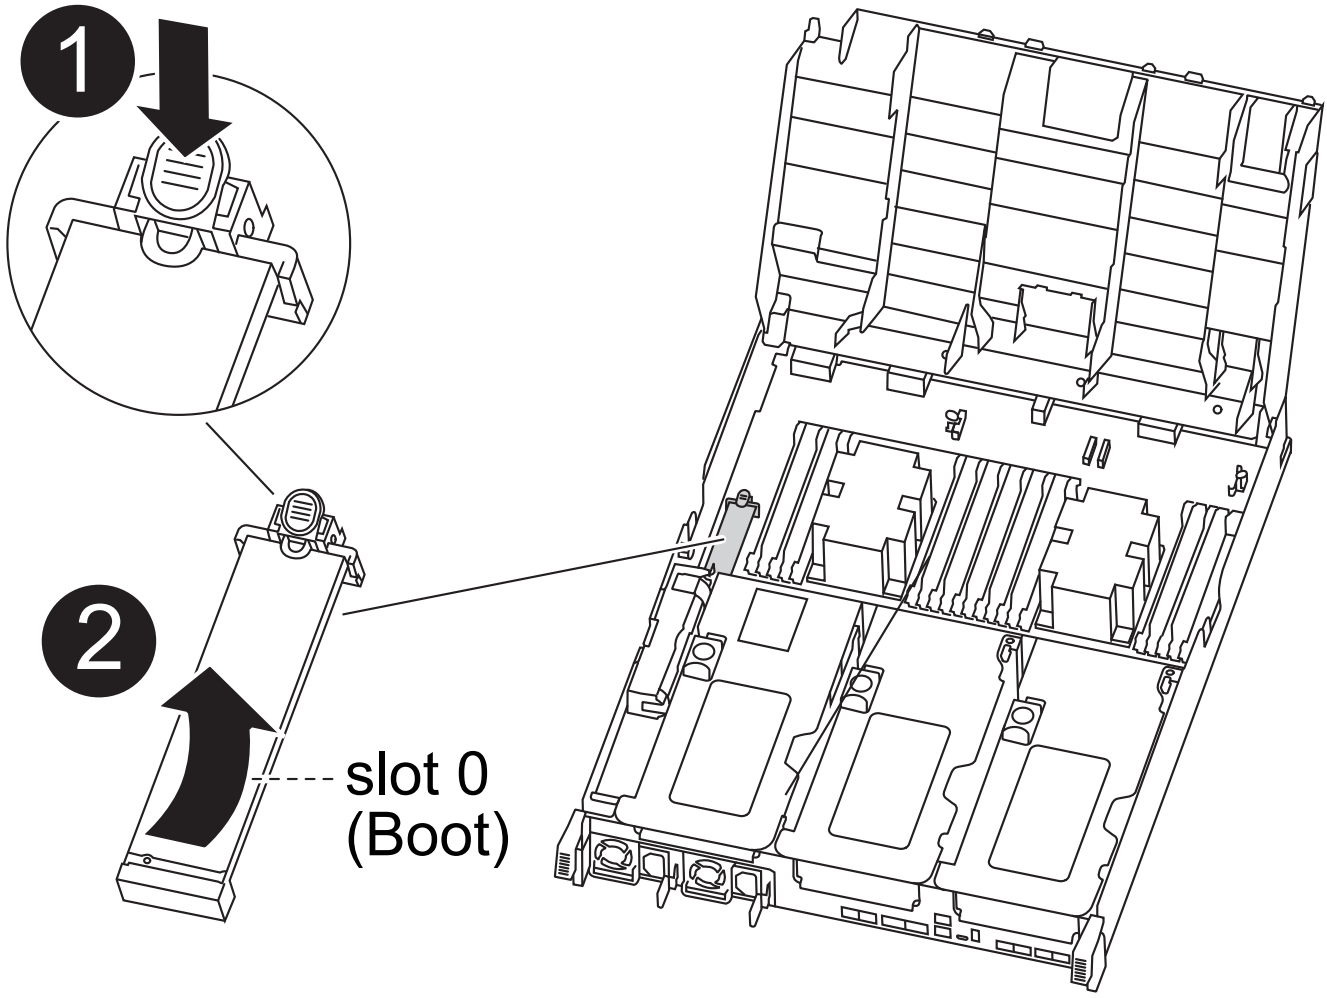

1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.

b. Rotate the boot media up and gently pull the boot media out of the socket.

10. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
11. Check the boot media to make sure that it is seated squarely and completely in the socket.

    If necessary, remove the boot media and reseat it into the socket.

12. Lock the boot media in place:
    a. Rotate the boot media down toward the motherboard.
    b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
    c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
13. Close the air duct.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - ASA C400

After installing the new boot media device in your ASA C400 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - ASA C400

If a component in your ASA C400 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - ASA C400

Get started with replacing the boot media in your ASA C400 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - ASA C400

Before replacing the boot media in your ASA C400 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - ASA C400

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - ASA C400

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster configuration

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 3: Controller is in a two-node Metrocluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes

that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the boot media and prepare for manual boot recovery - ASA C400

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

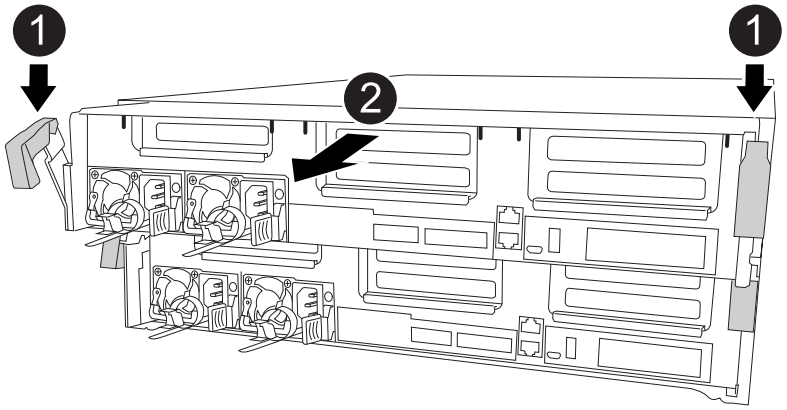

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 2: Replace the boot media

You must locate the boot media in the controller module (see the FRU map on the controller module), and then follow the directions to replace it.

Before you begin

Although the contents of the boot media is encrypted, it is a best practice to erase the contents of the boot media before replacing it. For more information, see the Statement of Volatility for your system on the NetApp Support Site.

> You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the boot media.

Animation - Replace the boot media

Steps

1. Open the air duct:

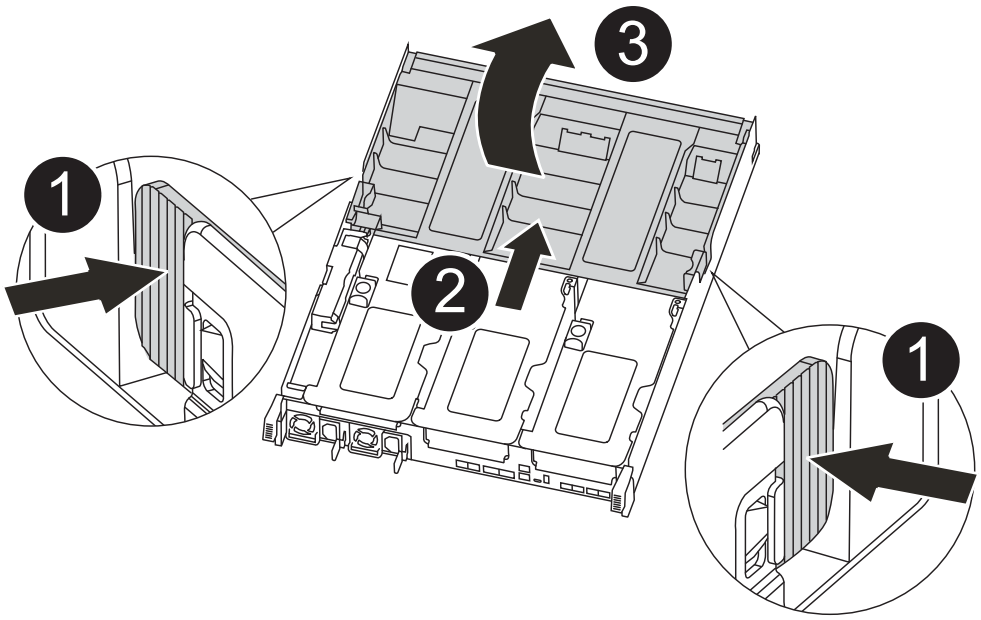

1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

2. Locate and remove the boot media from the controller module:

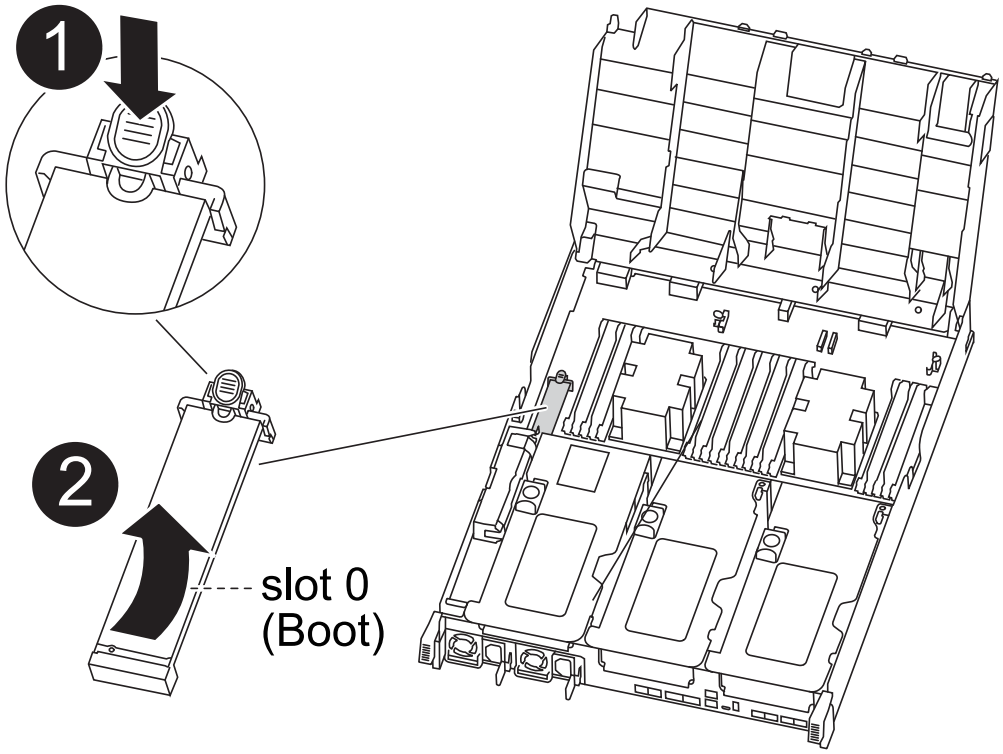

1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
b. Rotate the boot media up and gently pull the boot media out of the socket.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Lock the boot media in place:
   a. Rotate the boot media down toward the motherboard.
   b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
   c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
6. Close the air duct.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed does not have a boot image, so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the `var` file system.

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      There are two folders in the unzipped service image file:

      - `boot`
      - `efi`

   c. Copy the `efi` folder to the top directory on the USB flash drive.

      If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

      The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

   d. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.
5. Plug the power cable into the power supply and reinstall the power cable retainer.
6. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.
7. Complete the installation of the controller module:

a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

   The locking latches rise when the controller module is fully seated.

   > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

d. If you have not already done so, reinstall the cable management device.

8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then `halt` the controller to boot to LOADER.

9. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:

   a. Boot to Maintenance mode: `boot_ontap maint`

   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t` *`initiator adapter_name`*

   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - ASA C400

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `...would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore encryption - ASA C400

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    `security key-manager key query -restored false.`

    ```
    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    `storage failover giveback -fromnode local`

13. Restore automatic giveback, if you disabled it, by entering the following command.

    `storage failover modify -node local -auto-giveback true`

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    `system node autosupport invoke -node * -type all -message MAINT=END`

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*******************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*******************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - ASA C400

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - ASA C400

To replace the chassis, you must move the fans and controller modules from the impaired chassis to the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial

outage in a multinode cluster.

Shut down the controllers - ASA C400

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Shut down the controllers when replacing a chassis

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Shut down a controller in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the

`-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Move the fans, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move the fans

To move the fan modules to the replacement chassis when replacing the chassis, you must perform a specific sequence of tasks.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.
4. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.
5. Set the fan module aside.
6. Repeat the preceding steps for any remaining fan modules.

7. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.
10. Repeat these steps for the remaining fan modules.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
   b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - ASA C400

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for *HA-state* can be one of the following:

      - `ha`
      - `mcc`
      - `mcc-2n`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

4. Reinstall the bezel on the front of the system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - ASA C400

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement node* is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - ASA C400

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the controller module hardware - ASA C400

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

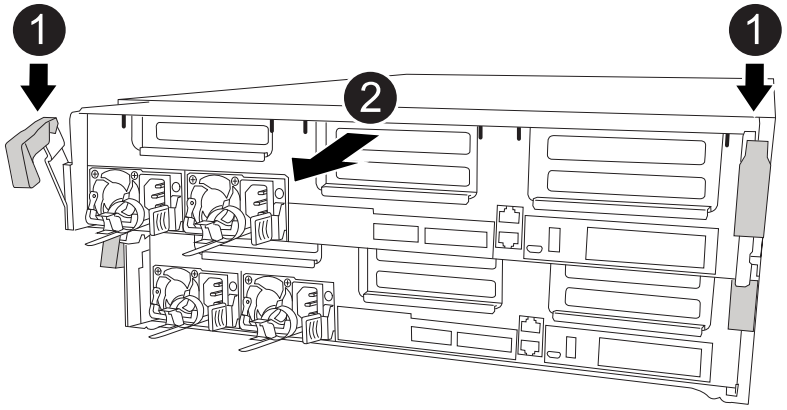

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. On the replacement controller module, open the air duct and remove the empty risers from the controller module using the animation, illustration, or the written steps:

Animation - Remove the empty risers from the replacement controller module

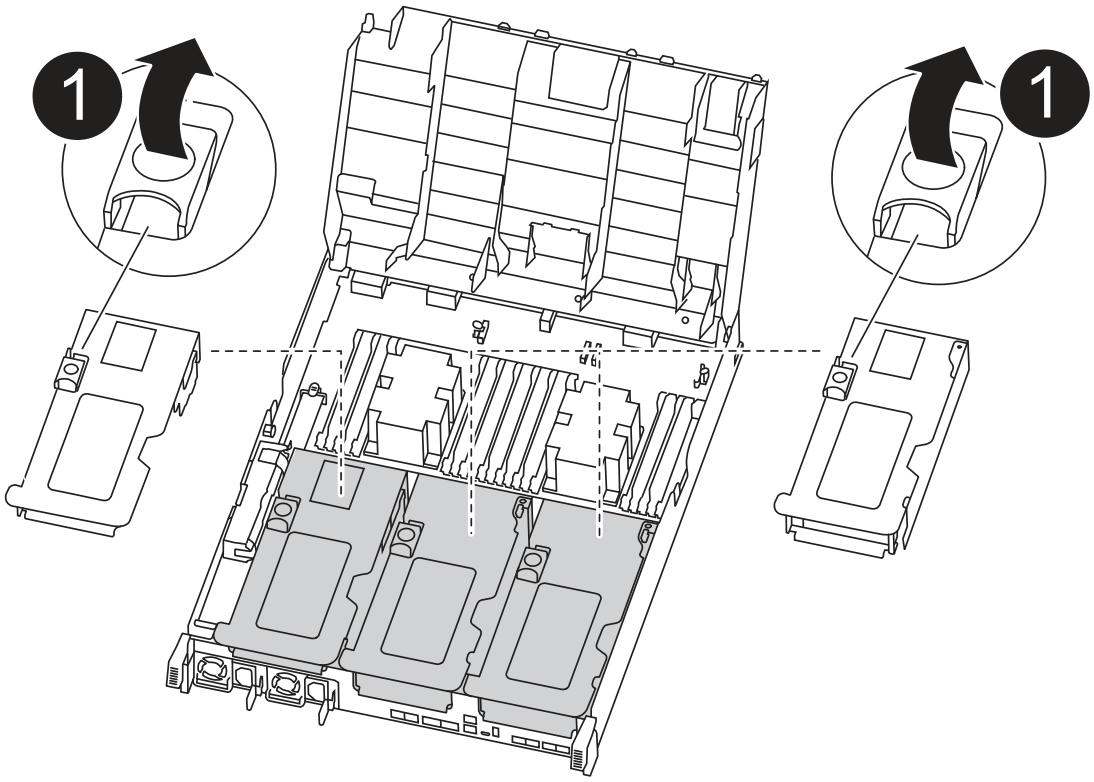

1	Riser latches

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
c. Rotate the riser locking latch on the left side of riser 1 up and toward air duct, lift the riser up, and then set it aside.
d. Repeat the previous step for the remaining risers.

Step 2: Move the power supplies

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

You can use the following animation, illustration, or the written steps to move the power supplies to the replacement controller module.

Animation - Move the power supplies

1. Remove the power supply:

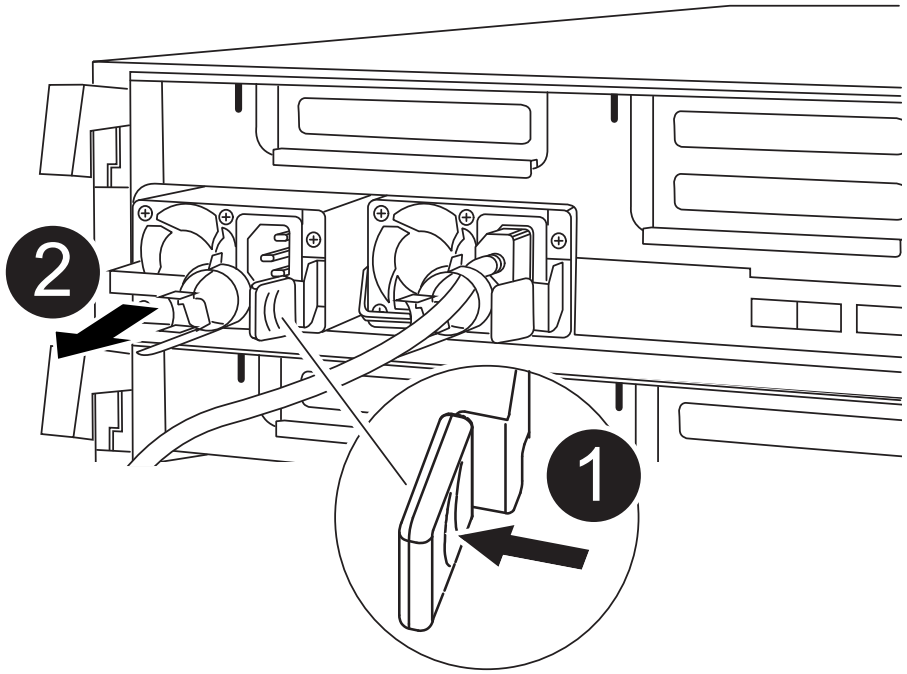

1	PSU locking tab
2	Power cable retainer

a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.

b. Press the blue locking tab to release the power supply from the chassis.

c. Using both hands, pull the power supply out of the chassis, and then set it aside.

1. Move the power supply to the new controller module, and then install it.
2. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

3. Repeat the preceding steps for any remaining power supplies.

Step 3: Move the NVDIMM battery

To move the NVDIMM battery from the impaired controller module to the replacement controller module, you must perform a specific sequence of steps.

You can use the following animation, illustration, or the written steps to move the NVDIMM battery from the impaired controller module to the replacement controller module.

Animation - Move the NVDIMM battery

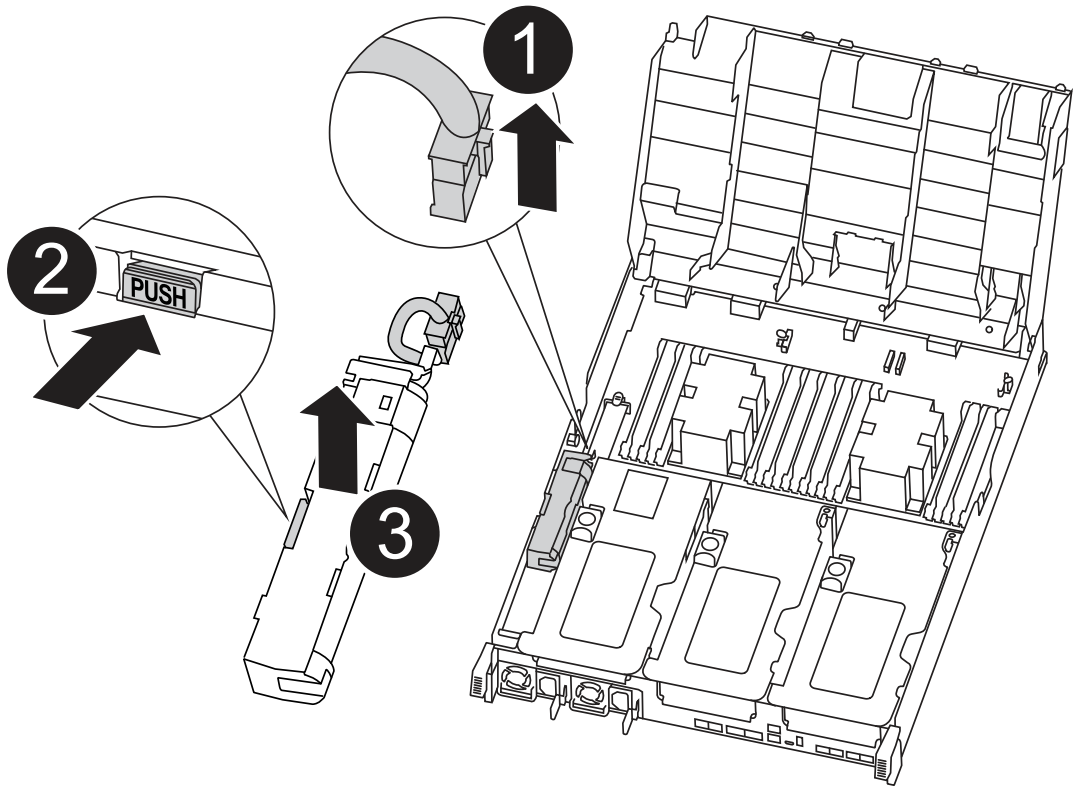

1	NVDIMM battery plug
2	NVDIMM battery locking tab
3	NVDIMM battery

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.
3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.

Do not plug the battery cable back into the motherboard until instructed to do so.

Step 4: Move the boot media

You must locate the boot media, and then follow the directions to remove it from the impaired controller module and insert it into the replacement controller module.

You can use the following animation, illustration, or the written steps to move the boot media from the impaired controller module to the replacement controller module.

Animation - Move the boot media

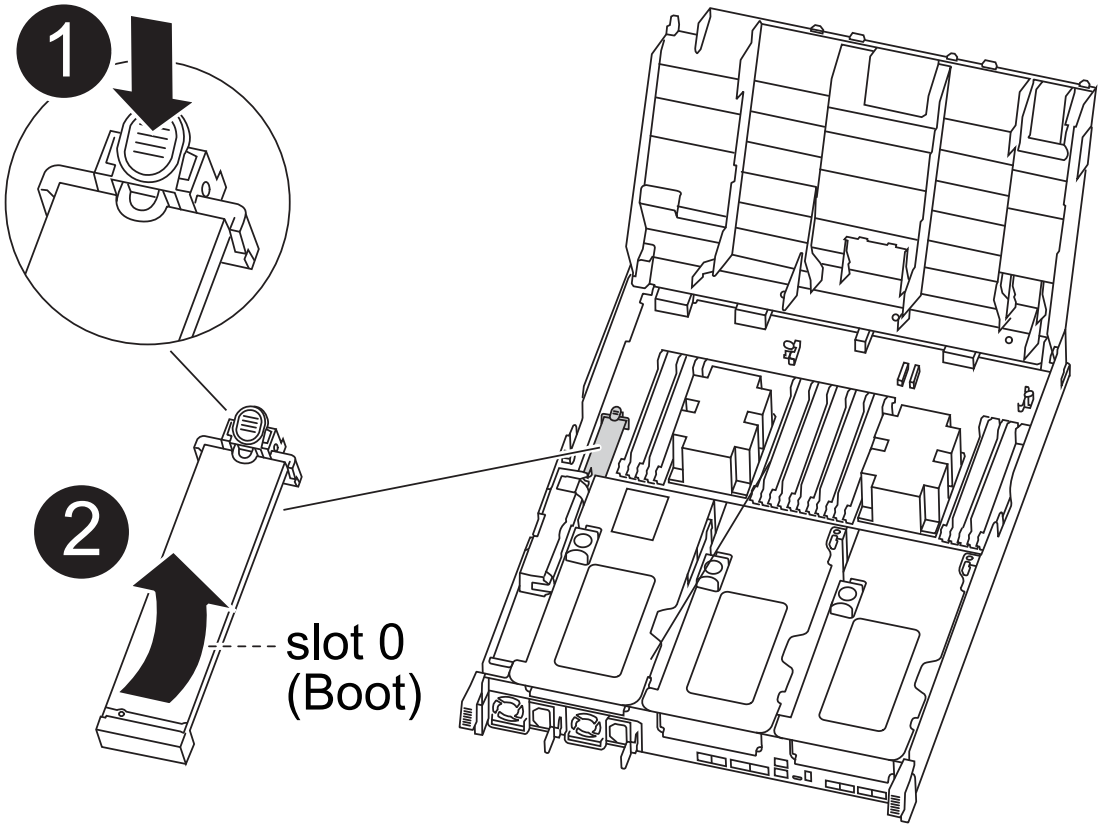

1	Boot media locking tab
2	Boot media

1. Locate and remove the boot media from the controller module:
   a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
   b. Rotate the boot media up and gently pull the boot media out of the socket.
2. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
3. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

4. Lock the boot media in place:

a. Rotate the boot media down toward the motherboard.

b. Press the blue locking button so that it is in the open position.

c. Placing your fingers at the end of the boot media by the blue button, firmly push down on the boot media end to engage the blue locking button.

Step 5: Move the PCIe risers and mezzanine card

As part of the controller replacement process, you must move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

You can use the following animations, illustrations, or the written steps to move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

Moving PCIe riser 1 and 2 (left and middle risers):

Animation - Move PCI risers 1 and 2

Moving the mezzanine card and riser 3 (right riser):

Animation - Move the mezzanine card and riser 3

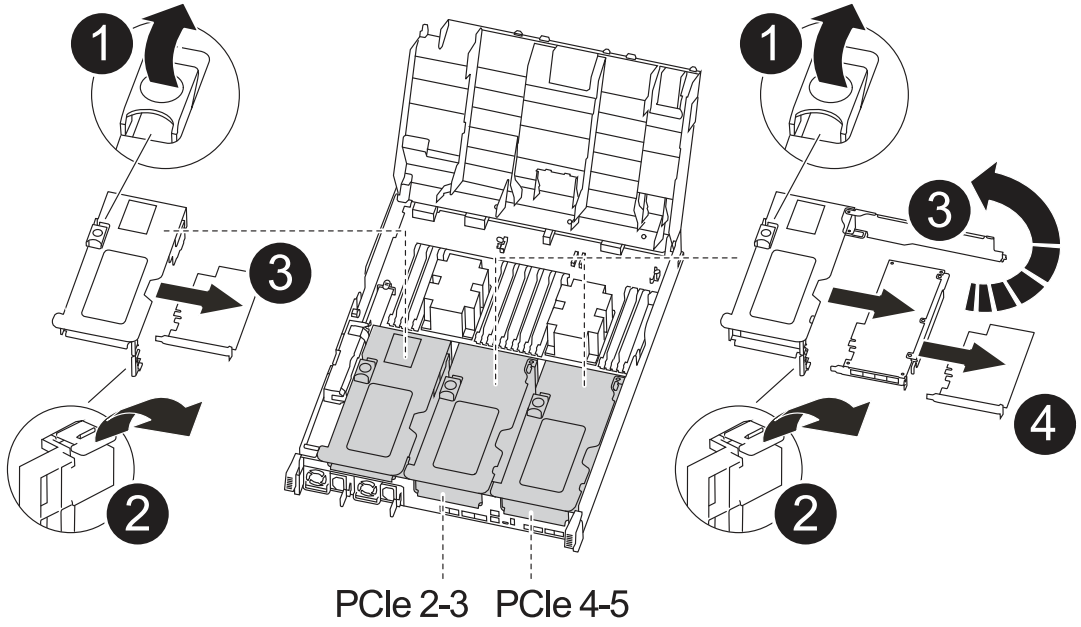

1	Riser locking latch
2	PCI card locking latch
3	PCI locking plate
4	PCI card

1. Move PCIe risers one and two from the impaired controller module to the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then move it to the replacement controller module.
   d. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins, push the riser squarely into the socket on the motherboard, and then rotate the latch down flush with the sheet metal on the riser.
   e. Repeat this step for riser number 2.
2. Remove riser number 3, remove the mezzanine card, and install both into the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then set it aside on a stable, flat surface.
   d. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket, and then move it to the replacement controller module.
   e. Install the mezzanine in the replacement controller and secure it with the thumbscrews.
   f. Install the third riser in the replacement controller module.

Step 6: Move the DIMMs

You need to locate the DIMMs, and then move them from the impaired controller module to the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

You can use the following animation, illustration, or the written steps to move the DIMMs from the impaired controller module to the replacement controller module.

Animation - Move the DIMMs

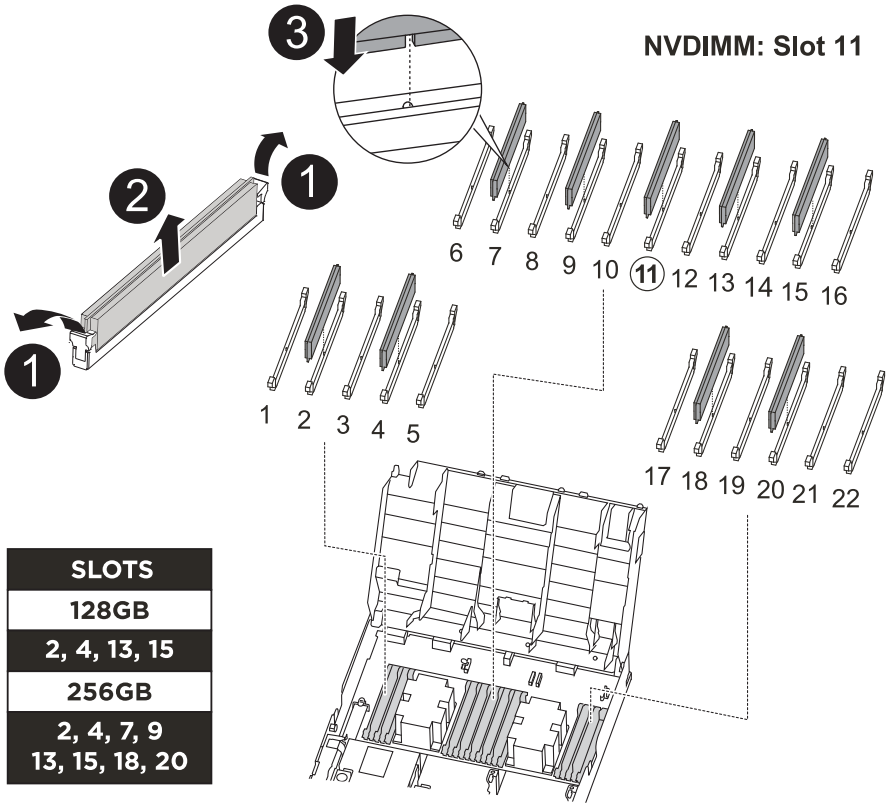

1	DIMM locking tabs
2	DIMM
3	DIMM socket

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Verify that the NVDIMM battery is not plugged into the new controller module.
4. Move the DIMMs from the impaired controller module to the replacement controller module:

   Make sure that you install the each DIMM into the same slot it occupied in the impaired controller module.

   a. Eject the DIMM from its slot by slowly pushing apart the DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

      Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   b. Locate the corresponding DIMM slot on the replacement controller module.

   c. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the

DIMM squarely into the socket.

The DIMMs fit tightly in the socket, but should go in easily. If not, realign the DIMM with the socket and reinsert it.

d. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

e. Repeat these substeps for the remaining DIMMs.

5. Plug the NVDIMM battery into the motherboard.

Make sure that the plug locks down onto the controller module.

Step 7: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

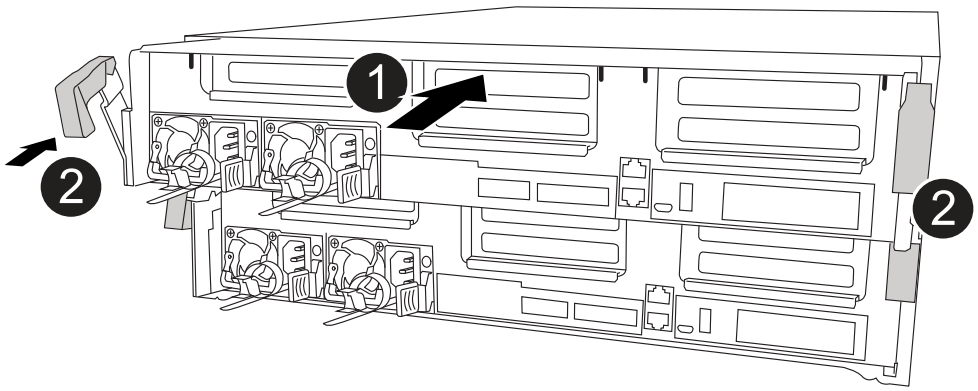

1	Slide controller into the chassis
2	Locking latches

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:

a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

Restore and verify the system configuration - ASA C400

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - ASA C400

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.

c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.

d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`

   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *`local-node-name`* `partner savecore -s`

d. Return to the admin privilege level: `set -privilege admin`

5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at

which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The *replacement* controller is the current owner of the disks on the disaster site.

     Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - ASA C400

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no

configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    ------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - ASA C400

You must replace a DIMM in the controller module when your system registers an increasing number of correctable error correction codes (ECC); failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

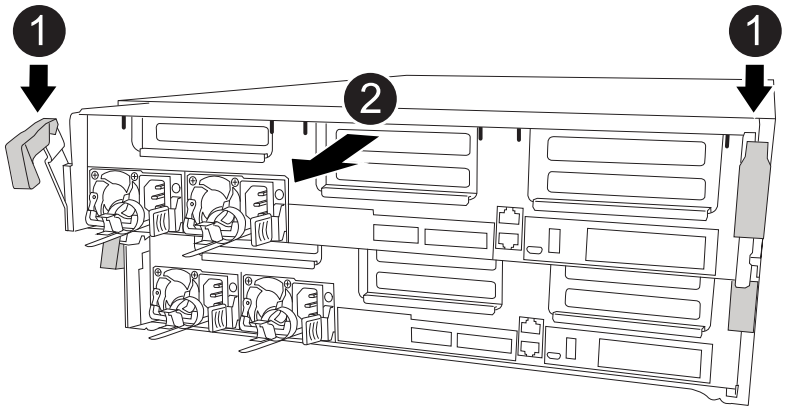

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace system DIMMs

Replacing a system DIMM involves identifying the target DIMM through the associated error message, locating the target DIMM using the FRU map on the air duct, and then replacing the DIMM.

You can use the following animation, illustration, or the written steps to replace a system DIMM.

> The animation and illustration show empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace a system DIMM

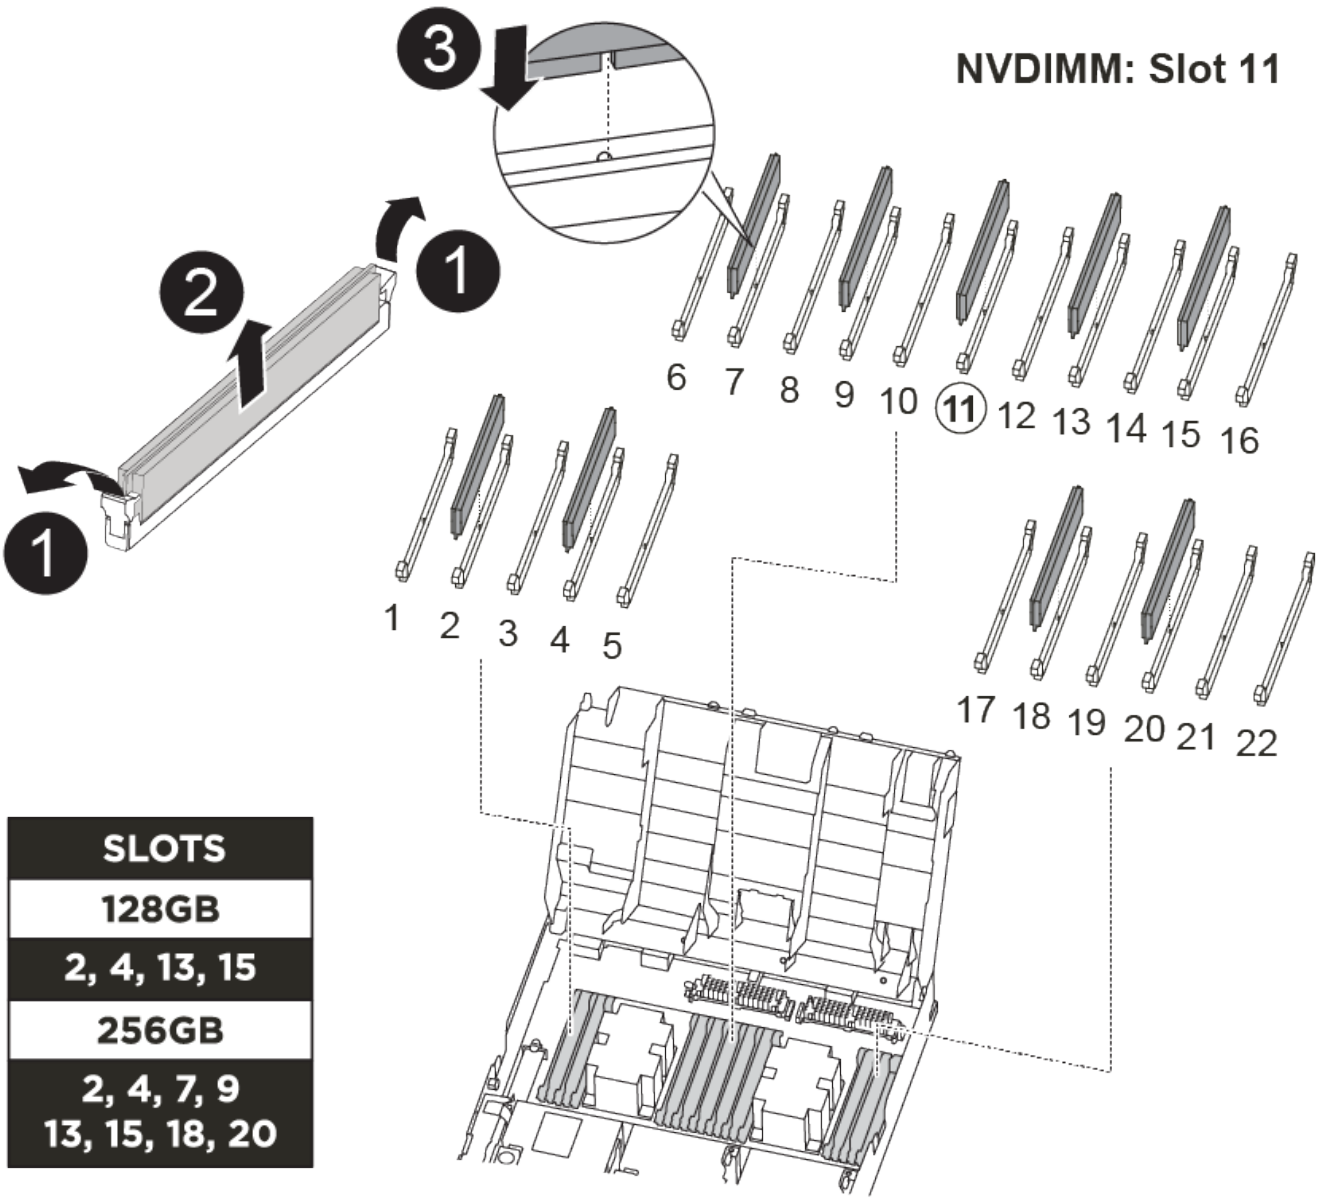

1	DIMM locking tabs
2	DIMM
3	DIMM socket

The DIMMs are located in sockets 2, 4, 13, and 15. The NVDIMM is located in slot 11.

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely

open position.

2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
4. Eject the DIMM from its socket by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the socket.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

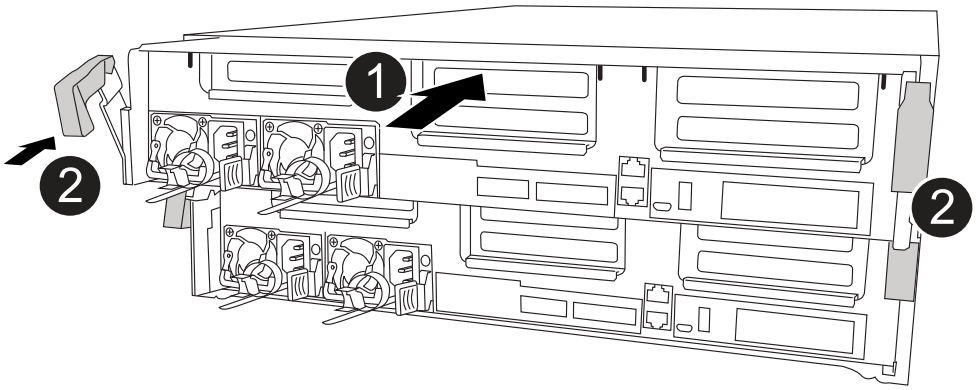

1	Controller module
2	Controller locking latches

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
   b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.
   g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto`

```
-giveback true
```

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    ------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a fan module - ASA C400

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

You can use the following animation, illustration, or the written steps to hot-swap a fan module.

Animation - Replace a fan

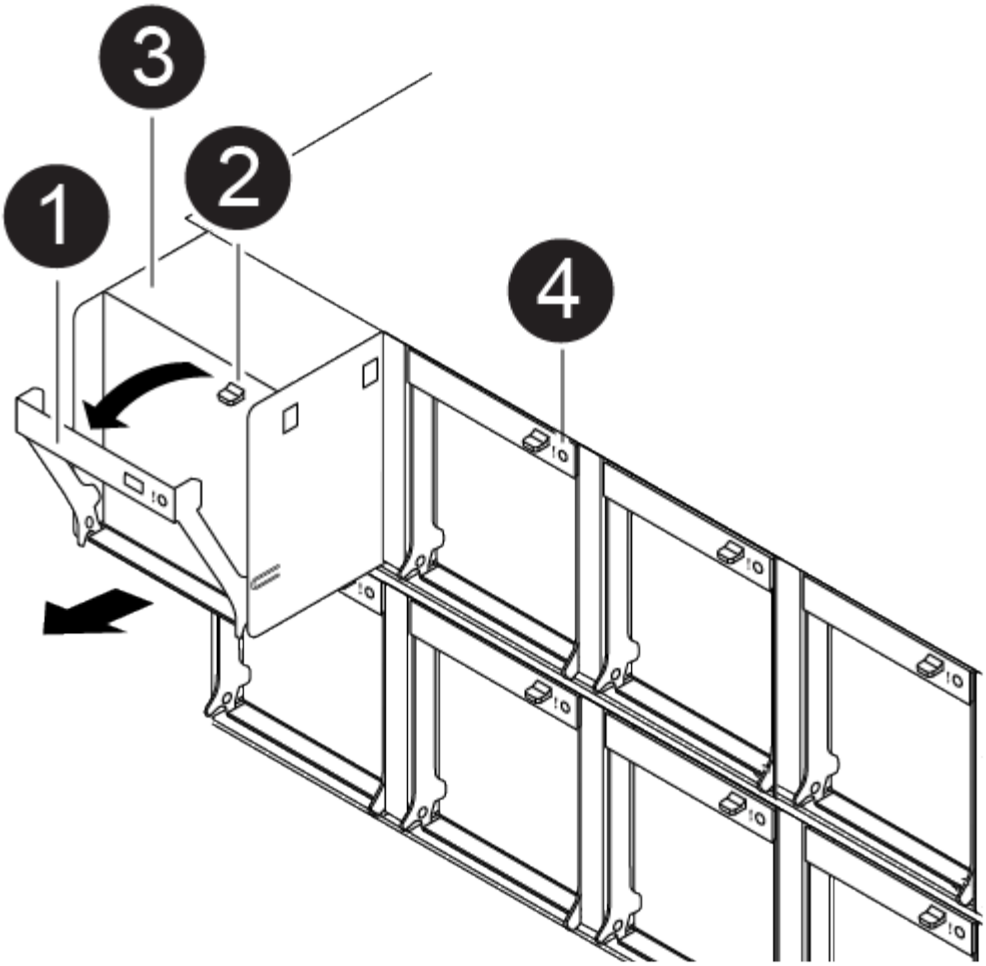

1	Fan handle
2	Locking tab
3	Fan
4	Status LED

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.

5. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

6. Set the fan module aside.
7. Insert the replacement fan module into the chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The Attention LED should not be lit after the fan is seated and has spun up to operational speed.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - ASA C400

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

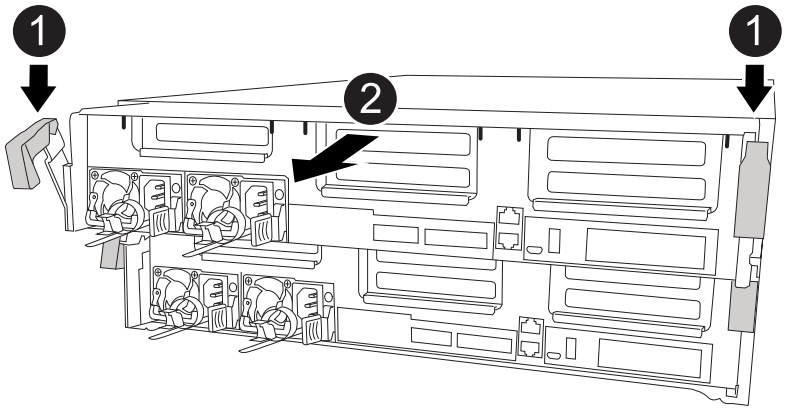

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module. See the FRU map inside the controller module to locate the NVDIMM battery.

The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.

You can use the following animation, illustration, or the written steps to replace the NVDIMM battery.

Animation - Replace the NVDIMM battery

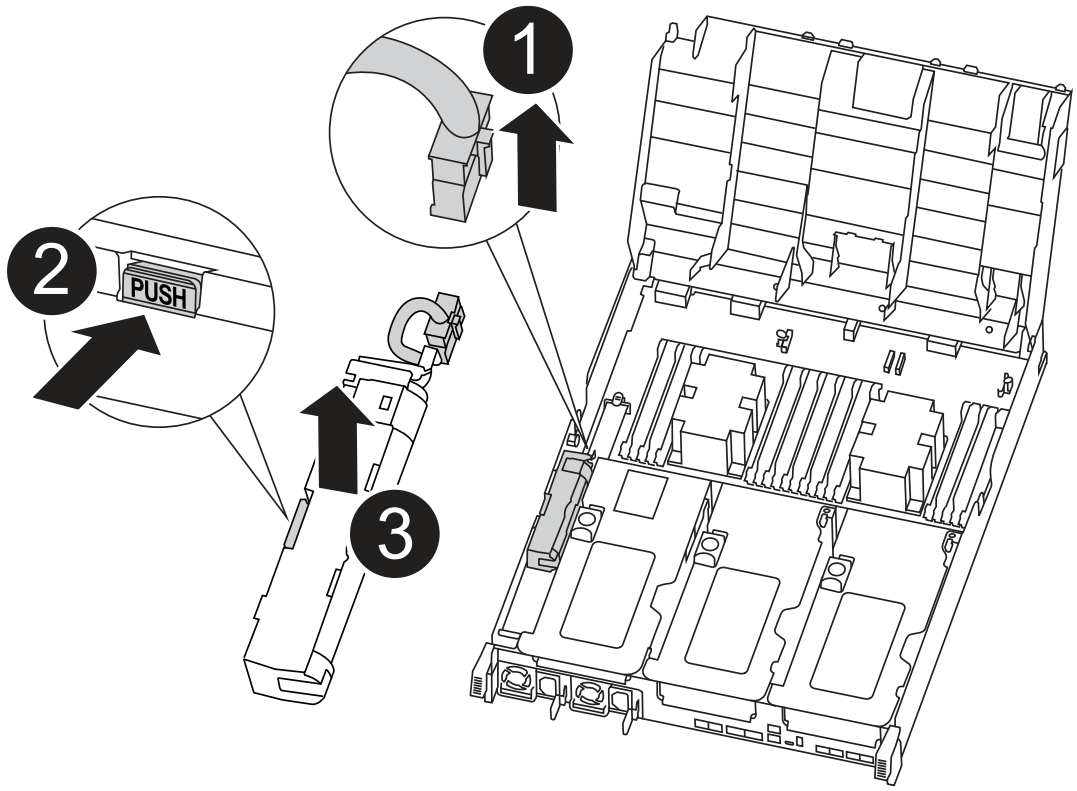

1	Battery plug
2	Locking tab
3	NVDIMM battery

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.
3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Remove the replacement battery from its package.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.
7. Plug the battery plug back into the controller module, and then close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

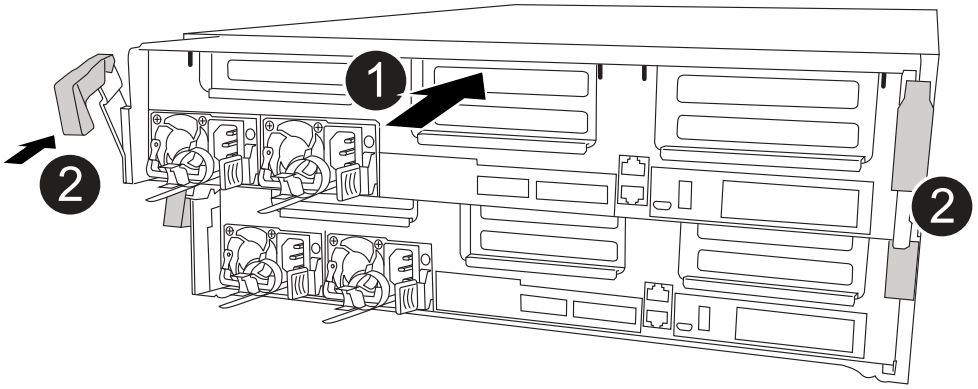

1	Controller module
2	Controller locking latches

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
   b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to

interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an NVDIMM - ASA C400

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

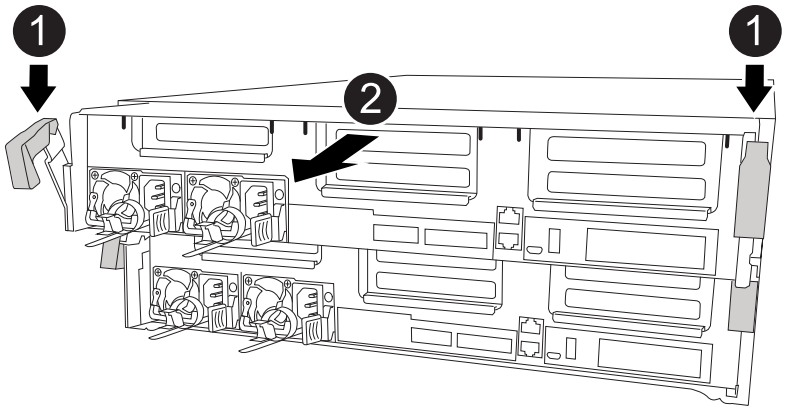

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the FRU map on top of the air duct or the FRU Map on the top of the slot 1 riser.

- The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.
- Although the contents of the NVDIMM is encrypted, it is a best practice to erase the contents of the NVDIMM before replacing it. For more information, see the Statement of Volatility on the NetApp Support

Site.

You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the NVDIMM.

The animation shows empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace the NVDIMM

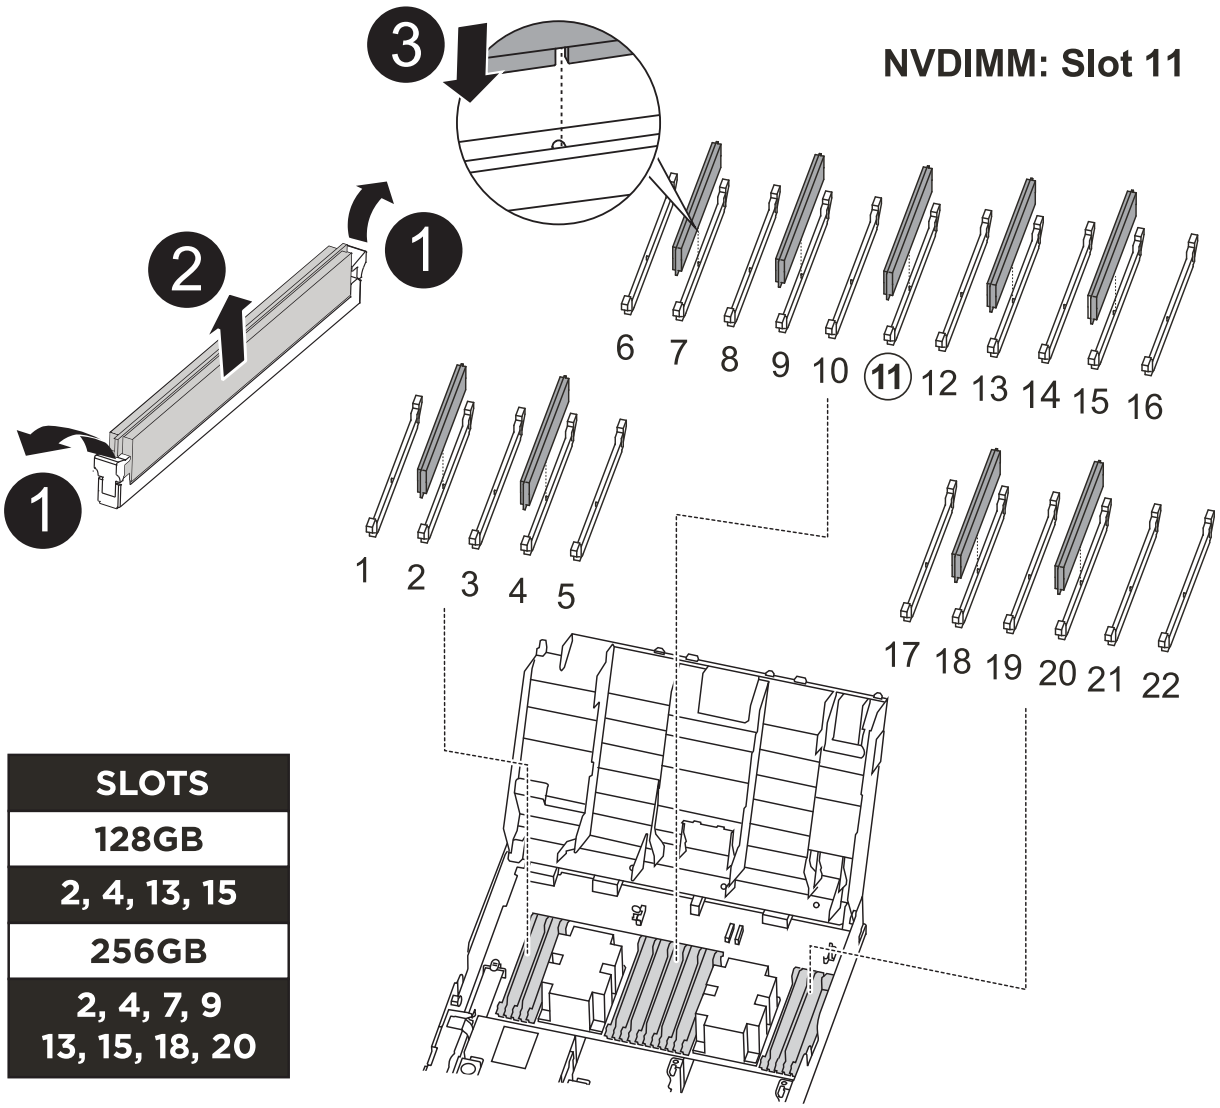

1	DIMM locking tabs
2	DIMM

3	DIMM socket

1. Open the air duct and then locate the NVDIMM in slot 11 on your controller module.

   The NVDIMM looks significantly different than system DIMMs.

2. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

3. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

4. Locate the slot where you are installing the NVDIMM.
5. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
7. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

> If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe or mezzanine card - ASA C400

To replace a PCIe or mezzanine card, you must disconnect the cables and any SFP and QSFP modules from the cards, replace the failed PCIe or mezzanine card, and then recable the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft

vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

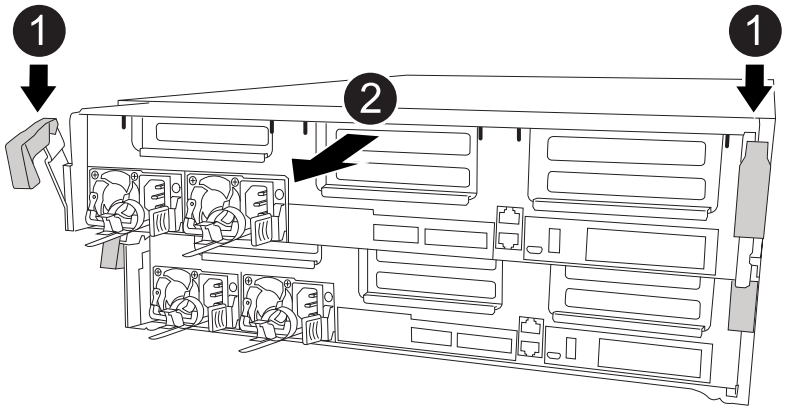

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace a PCIe card

To replace a PCIe card, you must locate the failed PCIe card, remove the riser that contains the card from the controller module, replace the card, and then reinstall the PCIe riser in the controller module.

You can use the following animation, illustration, or the written steps to replace a PCIe card.

Animation - Replace a PCIe card

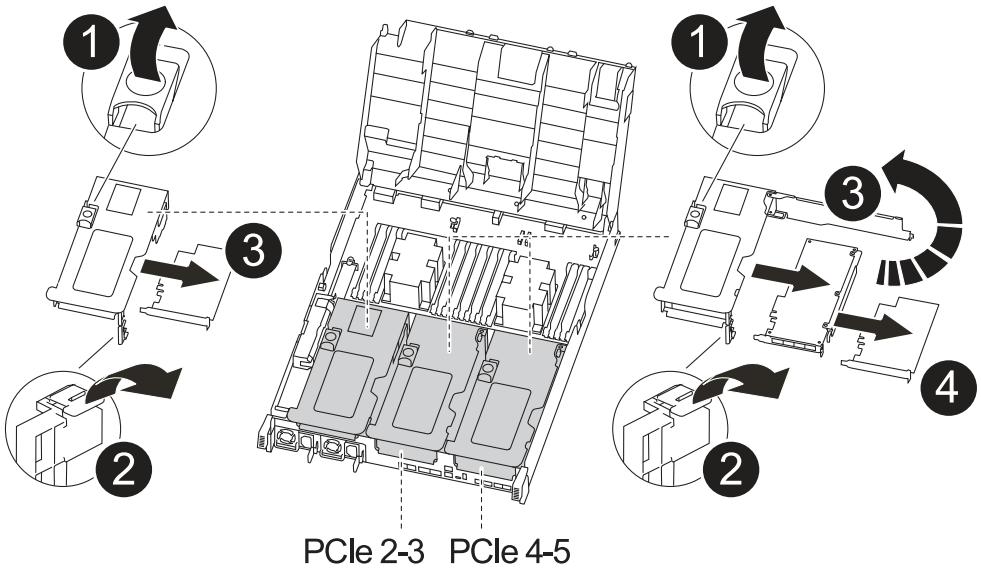

1	Riser locking latch
2	PCI card locking latch
3	PCI locking plate
4	PCI card

1. Remove the riser containing the card to be replaced:
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up straight up and set it aside on a stable flat surface,
2. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. For risers 2 and 3 only, swing the side panel up.
   d. Remove the PCIe card from the riser by gently pushing up on the bracket and lift the card straight out of the socket.
3. Install the replacement PCIe card in the riser by aligning the card with the socket, press the card into the socket and then close the side panel on the riser, if present.

Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The PCIe card must be fully and evenly seated in the slot.

If you are installing a card in the bottom slot and cannot see the card socket well, remove the top card so that you can see the card socket, install the card, and then reinstall the card you removed from the top slot.

4. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 4: Replace the mezzanine card

The mezzanine card is located under riser number 3 (slots 4 and 5). You must remove that riser to access the mezzanine card, replace the mezzanine card, and then reinstall riser number 3. See the FRU map on the controller module for more information.

You can use the following animation, illustration, or the written steps to replace the mezzanine card.

Animation - Replace the mezzanine card

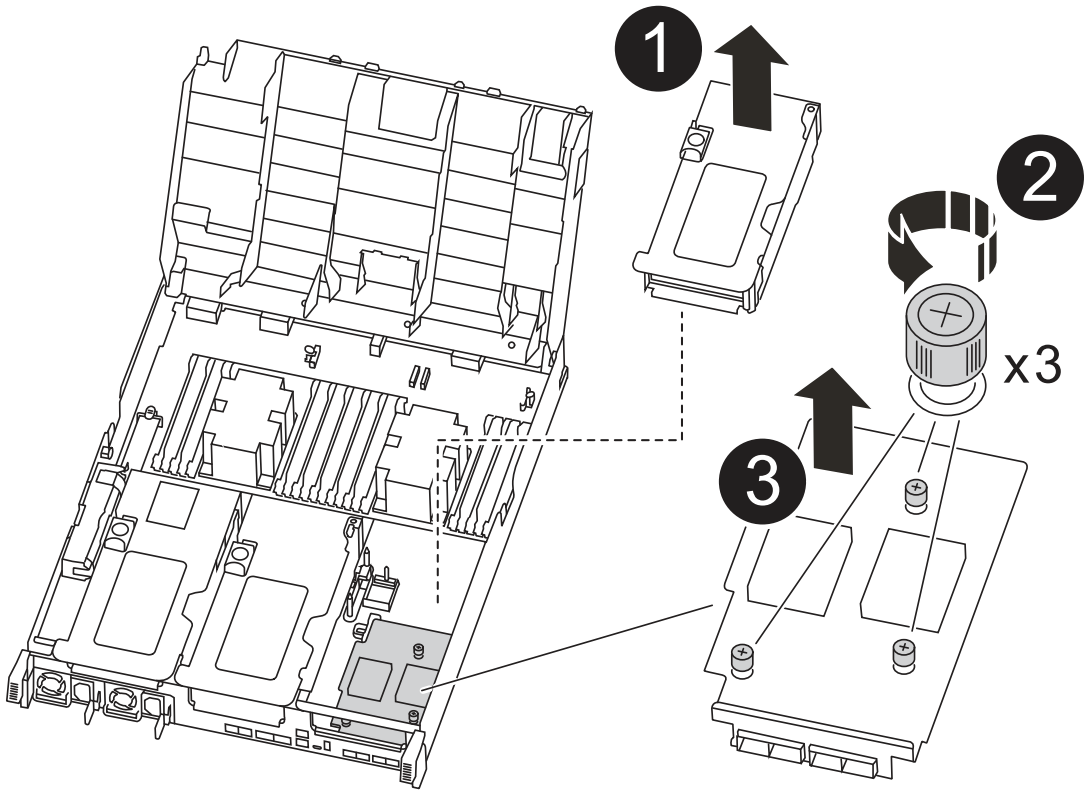

1	PCI riser

2	Riser thumbscrew
3	Riser card

1. Remove riser number 3 (slots 4 and 5):
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up, and then set it aside on a stable, flat surface.
2. Replace the mezzanine card:
   a. Remove any QSFP or SFP modules from the card.
   b. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket and set it aside.
   c. Align the replacement mezzanine card over the socket and the guide pins and gently push the card into the socket.
   d. Tighten the thumbscrews on the mezzanine card.
3. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 5: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the installation of the controller module:
   a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.

b. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

   The locking latches rise when the controller module is fully seated.

   > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

   > If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

5. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
6. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Restore the controller module to operation

To restore the controller, you must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 7: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 8: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replacing a power supply - ASA C400

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting the replacement PSU to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

You can use the following illustration with the written steps to replace the power supply.

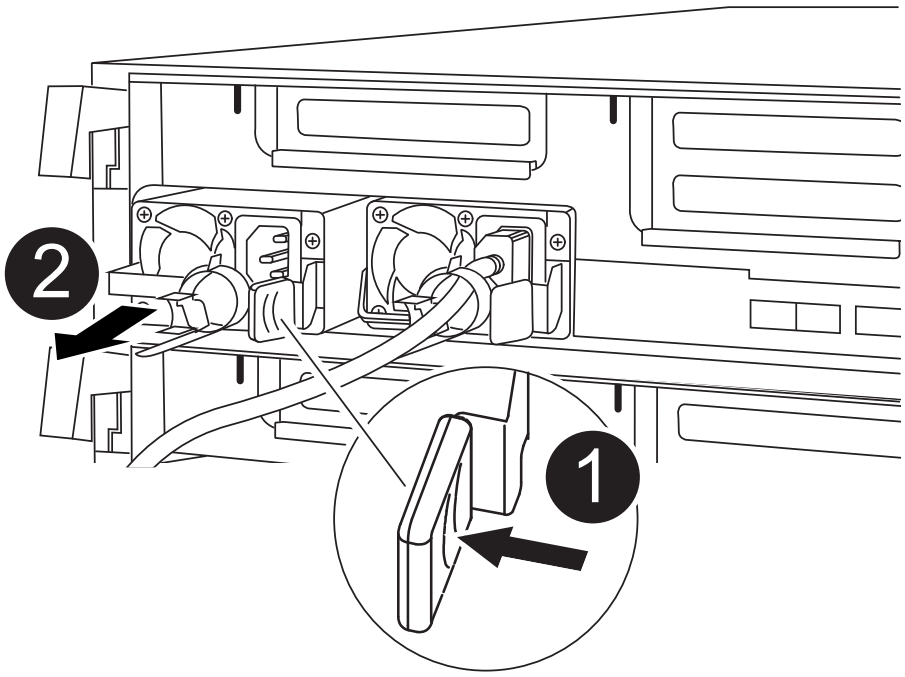

1	PSU locking tab
2	Power cable retainer

1. If you are not already grounded, properly ground yourself.

2. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
3. Disconnect the power supply:
   a. Open the power cable retainer, and then unplug the power cable from the power supply.
   b. Unplug the power cable from the power source.
4. Remove the power supply:
   a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
   b. Press the blue locking tab to release the power supply from the chassis.
   c. Using both hands, pull the power supply out of the chassis, and then set it aside.
5. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

6. Rotate the cam handle so that it is flush against the power supply.
7. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA C400

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

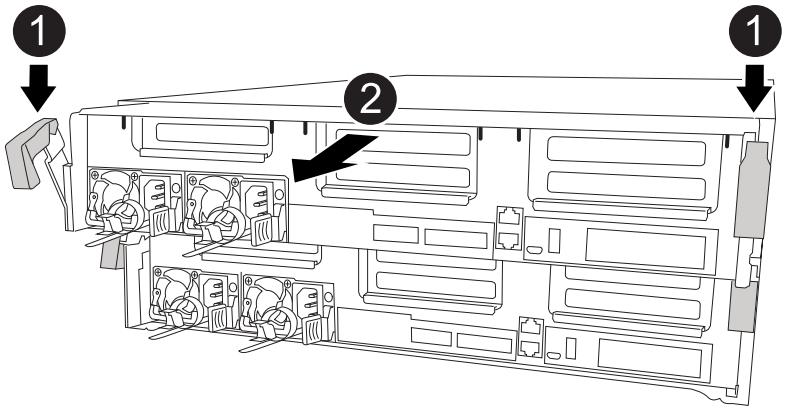

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Place the controller module on a stable, flat surface.

Step 3: Replace the RTC battery

You need to locate the RTC battery inside the controller module, and then follow the specific sequence of steps. See the FRU map inside the controller module for the location of the RTC battery.

You can use the following animation, illustration, or the written steps to replace the RTC battery.

Animation- Replace the RTC battery

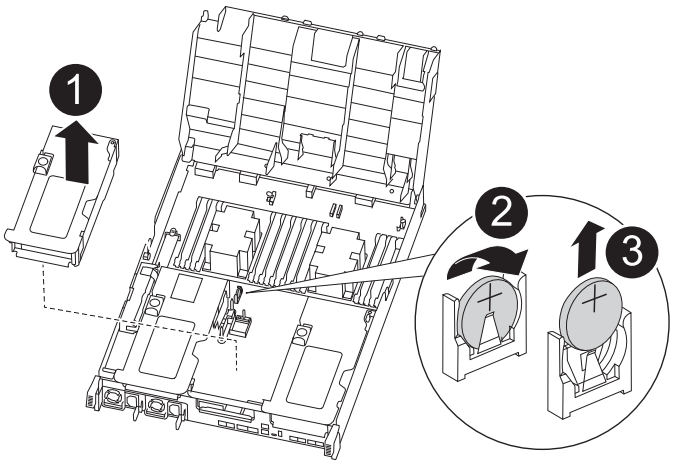

1	Middle riser
2	Remove RTC battery
3	Seat RTC battery

1. If you are not already grounded, properly ground yourself.
2. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
3. Locate, remove, and then replace the RTC battery:
   a. Using the FRU map, locate the RTC battery on the controller module.
   b. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

   > Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

   c. Remove the replacement battery from the antistatic shipping bag.
   d. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
5. Close the air duct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. If you have not already done so, reinstall the cable management device.
   d. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      > If your system stops at the boot menu, select the option to boot to LOADER.

6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA C800 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations (including ASA configurations), you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick steps - ASA C800

Quick start gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up.

Use the ASA C800 Installation and Setup Instructions if you are familiar with installing NetApp systems.

Videos - ASA C800

There are two videos - one showing how to rack and cable your system and one showing an example of using the System Manager Guided Setup to perform initial system configuration.

Video one of two: Hardware installation and cabling

The following video shows how to install and cable your new system.

"Animation - Installation and Setup of an ASA C800

Video two of two: Perform end-to-end software configuration

The following video shows end-to-end software configuration for systems running ONTAP 9.2 and later.

[] | *https://img.youtube.com/vi/Q6orVMyj94A?/maxresdefault.jpg*

Detailed steps - ASA C800

This section gives detailed step-by-step instructions for installing an ASA C800 system.

Step 1: Prepare for installation

To install your ASA C800 system, you need to create an account and register the system. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the NetApp Hardware Universe (HWU) for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

What you need

You need to provide the following at your site:

- Rack space for the storage system
  - 4U in an HA configuration for the platform
  - 2U for each NS224 storage shelf
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

Steps

1. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register (NetApp Product Registration) your system.
2. Download and install NetApp Downloads: Config Advisor on your laptop.
3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Connector type	Part number and length	Type of cable…	For…
100 GbE cable	X66211A-05 (112-00595), 0.5m X66211-1 (112-00573), 1m X66211-2 (112-00574), 2m X66211-5 (112-00576), 5m		• HA interconnect • Cluster interconnect network • Storage, Data
10 GbE cable	X6566B-3-R6 (112-00300), 3m; X6566B-5-R6 (112-00301), 5m		• Data
25 GbE cable	X66240A-2 (112-00598), 2m; X66240A-5 (112-00600), 5m		• Data
RJ-45 (order dependent)	Not applicable		• Management
Fibre Channel	X66250-2 (112-00342) 2m; X66250-5 (112-00344) 5m; X66250-15 (112-00346) 15m; X66250-30 (112-00347) 30m		• Network
Micro-USB console cable	Not applicable		• Console connection during software setup
Power cables	Not applicable		Connecting the PSUs to power source

4. Download and complete the Cluster Configuration Worksheet.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.

   Installing SuperRail into a four-post rack

2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

3. Attach cable management devices (as shown).

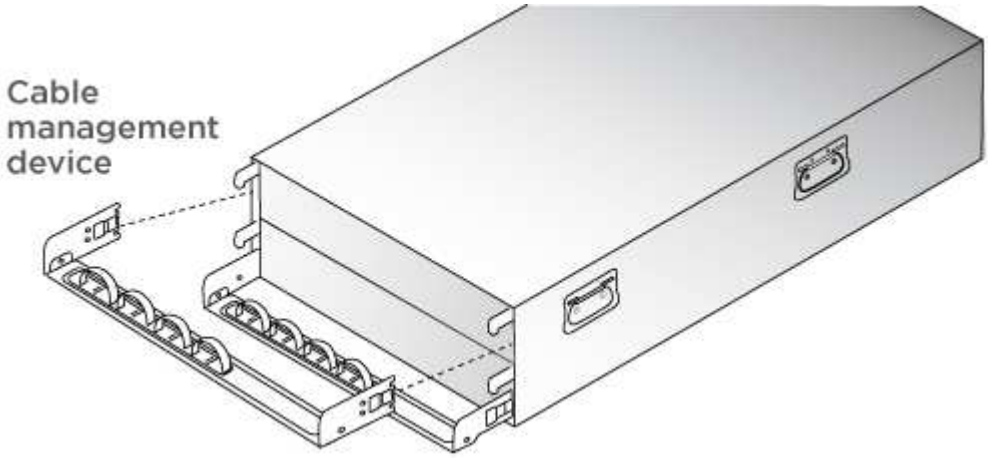

4. Place the bezel on the front of the system.

Step 3: Cable controllers

There is required cabling for your platform's cluster using the two-node switchless cluster method or the cluster interconnect network method. There is optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cable to a host network and storage.

Required cabling: Cable controllers to a cluster

Cable the controllers to a cluster by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Cable a two-node switchless cluster

Management network ports on the controllers are connected to switches. The HA interconnect and cluster interconnect ports are cabled on both controllers.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

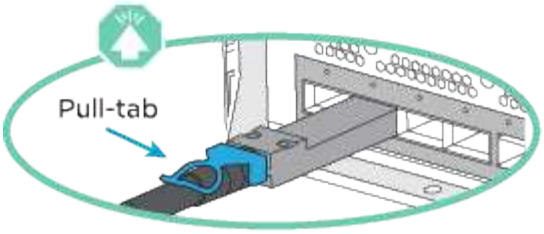

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or the tabulated steps to complete the cabling between the controllers and the switches:

Animation - Cable a two-node switchless cluster

Step	Perform on each controller module
1	Cable the HA interconnect ports: • e0b to e0b • e1b to e1b 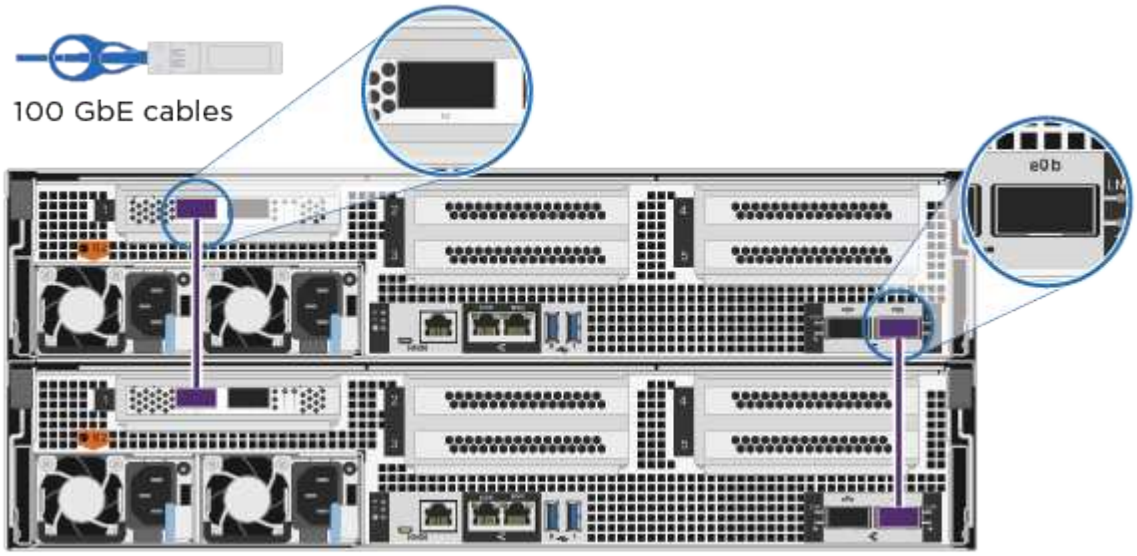

Step	Perform on each controller module
2	Cable the cluster interconnect ports: • e0a to e0a • e1a to e1a 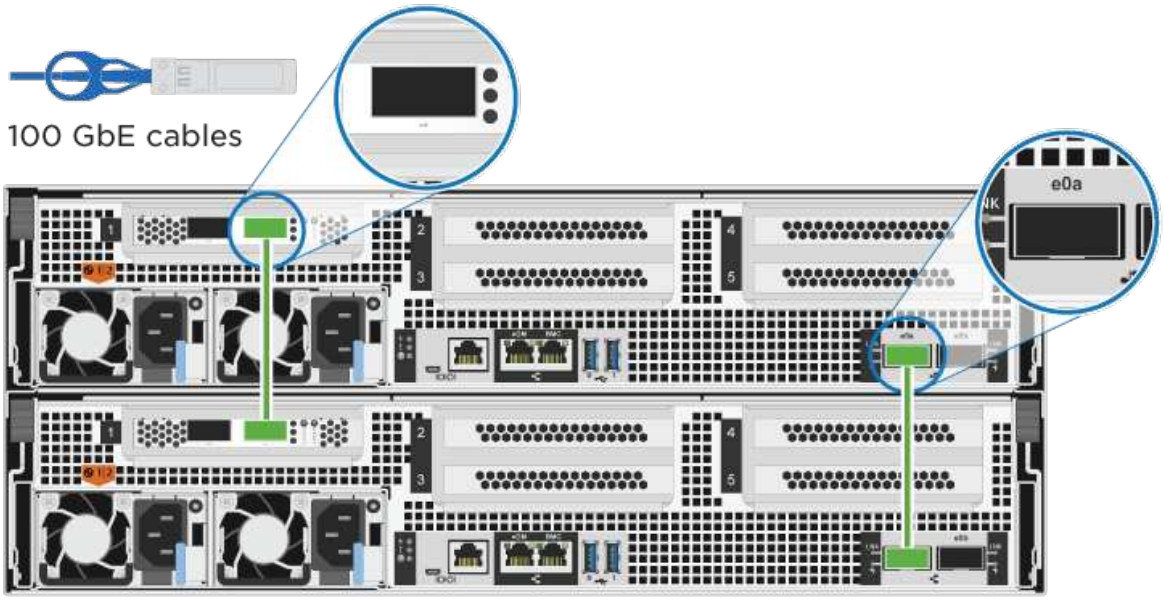
3	Cable the management ports to the management network switches 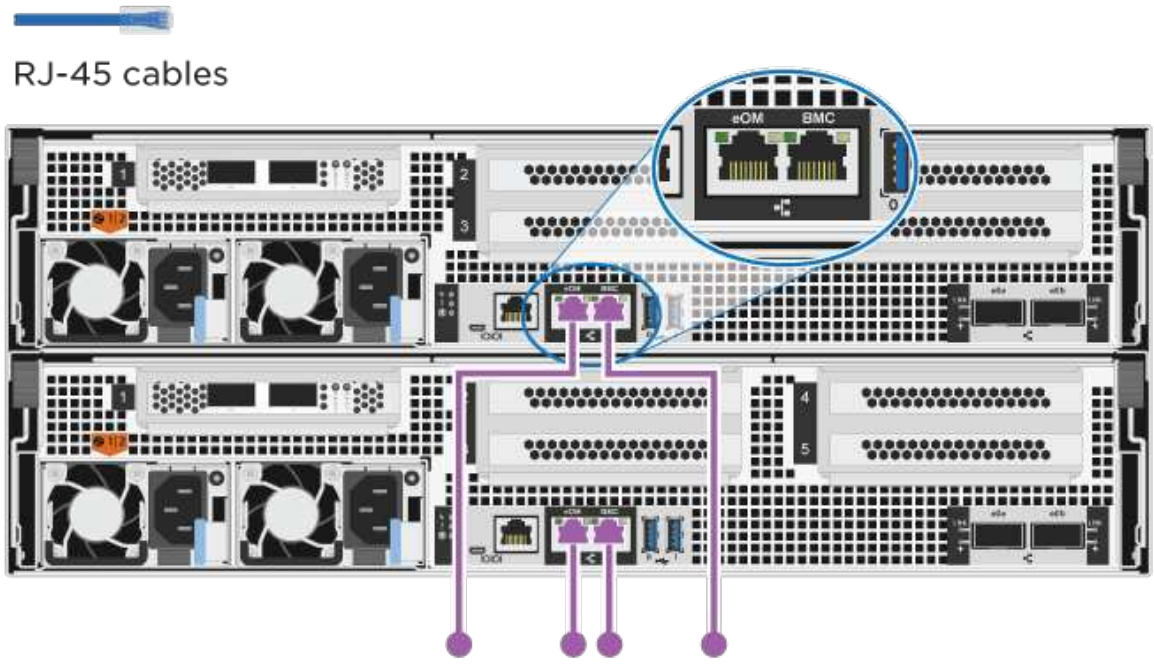
!	DO NOT plug in the power cords at this point.

2. To perform optional cabling, see:
   - Option 1: Cable to a Fibre Channel host network
   - Option 2: Cable to a 10GbE host network
   - Option 3: Cable the controllers to a single drive shelf
   - Option 4: Cable the controllers to two drive shelves
3. To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 2: Cable a switched cluster

Cluster interconnect and management network ports on the controllers are connected to switches while the HA interconnect ports are cabled on both controllers.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

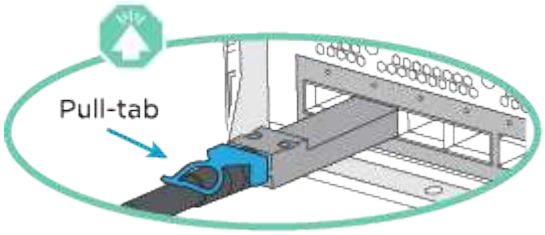

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or the tabulated steps to complete the cabling between the controllers and the switches:

   Animation - Cable a switched cluster

Step	Perform on each controller module
1	Cable the HA interconnect ports: • e0b to e0b • e1b to e1b 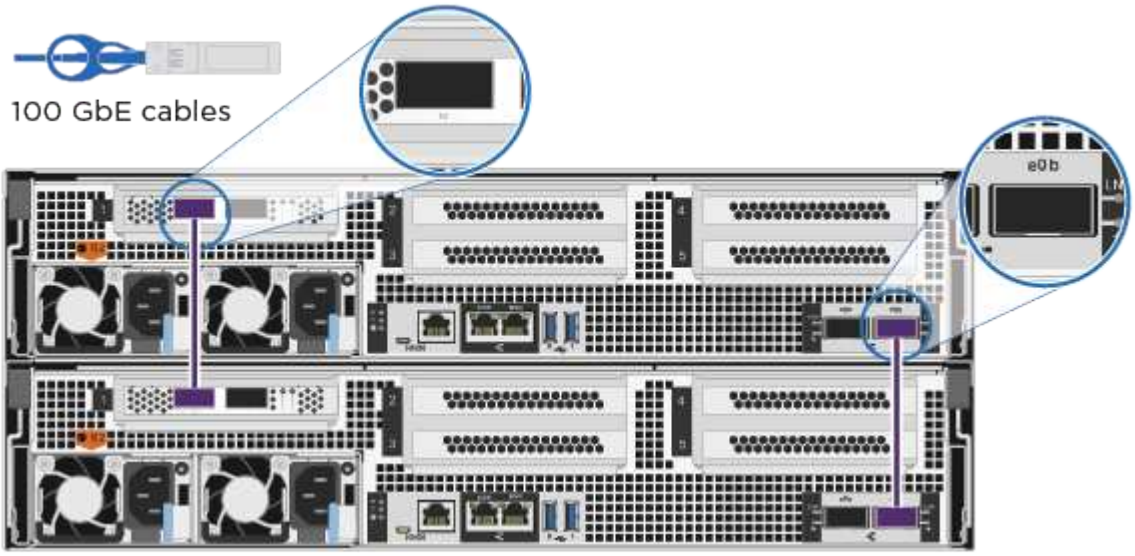
2	Cable the cluster interconnect ports to the 100 GbE cluster interconnect switches. • e0a • e1a 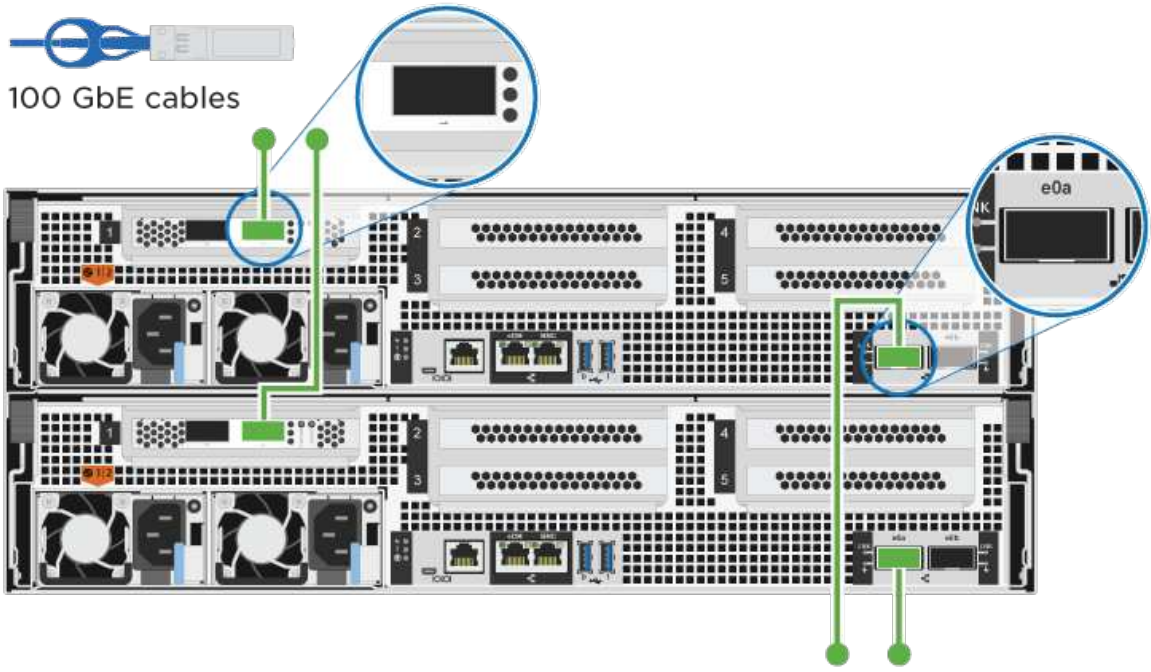

Step	Perform on each controller module
3	Cable the management ports to the management network switches 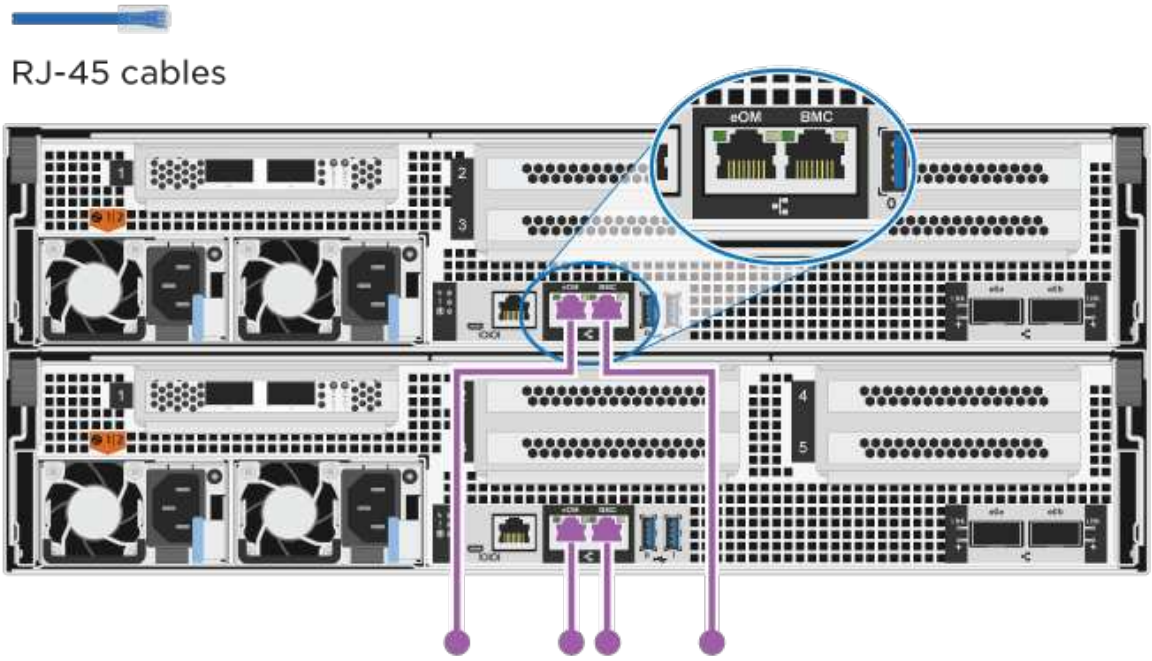
	DO NOT plug in the power cords at this point.

2. To perform optional cabling, see:
   - Option 1: Cable to a Fibre Channel host network
   - Option 2: Cable to a 10GbE host network
   - Option 3: Cable the controllers to a single drive shelf
   - Option 4: Cable the controllers to two drive shelves
3. To complete setting up your system, see Step 4: Complete system setup and configuration.

Optional cabling: Cable configuration-dependent options

You have configuration-dependent optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cabling to a host network and storage.

Option 1: Cable to a Fibre Channel host network

Fibre Channel ports on the controllers are connected to Fibre Channel host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

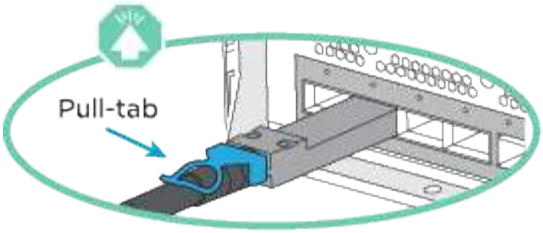

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Step	Perform on each controller module
1	Cable ports 2a through 2d to the FC host switches. 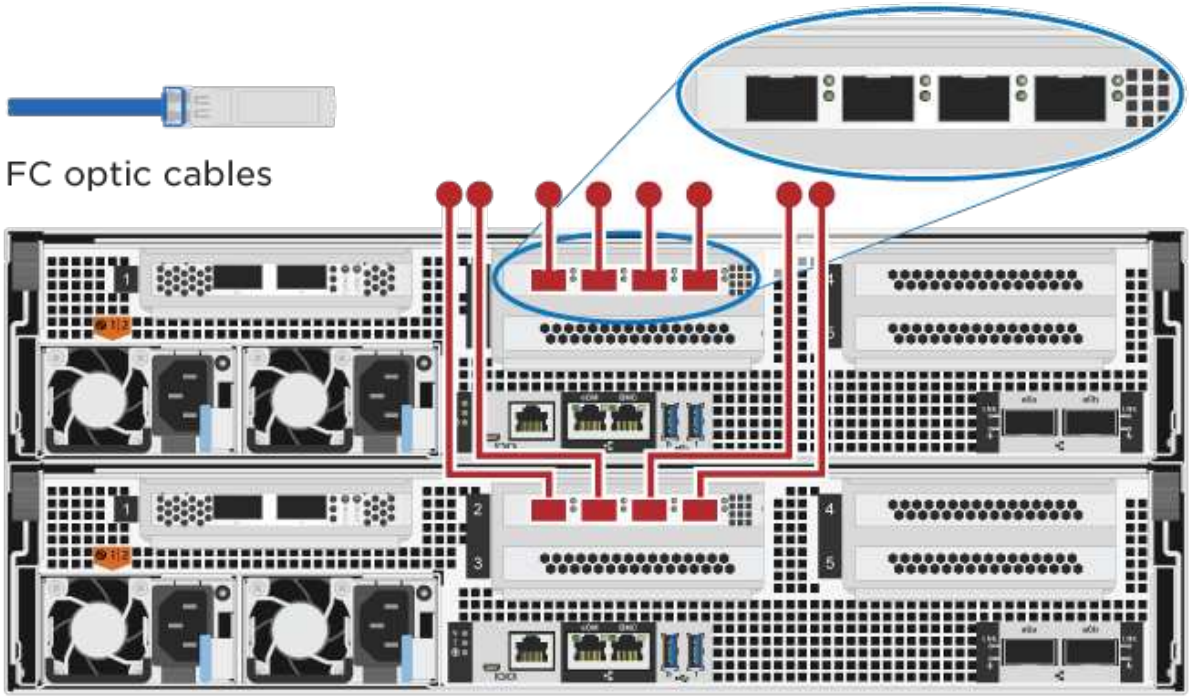
2	To perform other optional cabling, choose from: • Option 3: Cable the controllers to a single drive shelf • Option 4: Cable the controllers to two drive shelves
3	To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 2: Cable to a 10GbE host network

10GbE ports on the controllers are connected to 10GbE host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

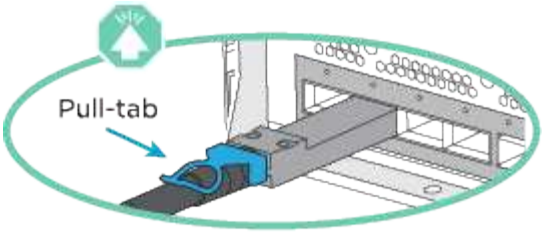

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Step	Perform on each controller module
1	Cable ports e4a through e4d to the 10GbE host network switches. 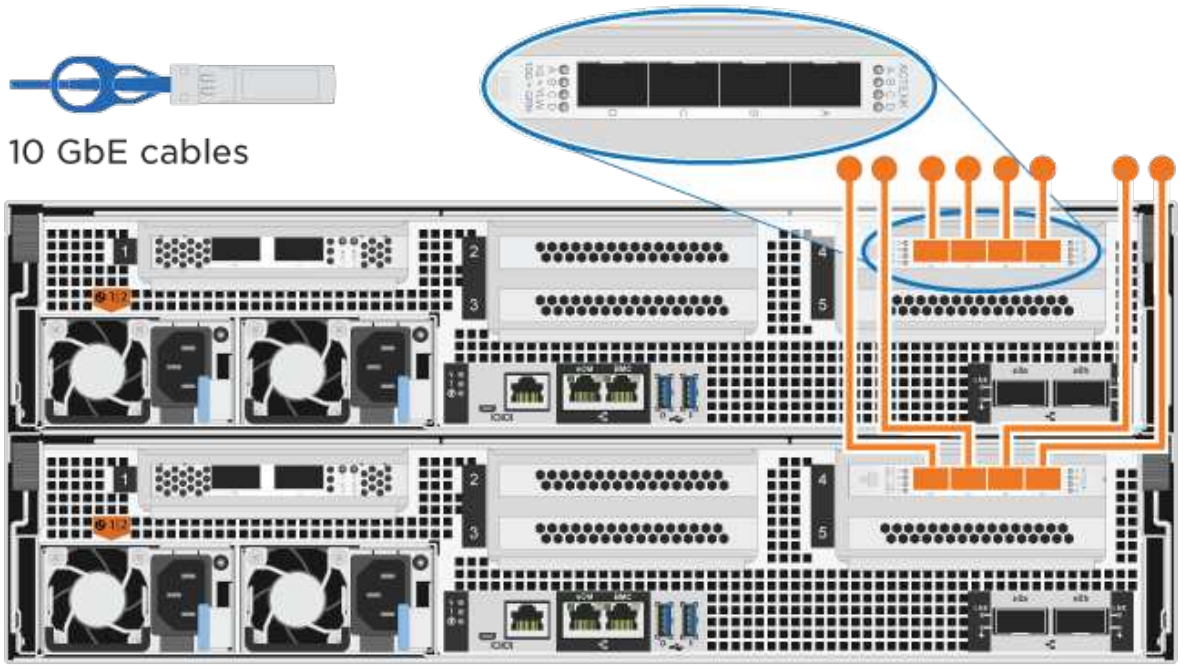
2	To perform other optional cabling, choose from: • Option 3: Cable the controllers to a single drive shelf • Option 4: Cable the controllers to two drive shelves
3	To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 3: Cable the controllers to a single drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

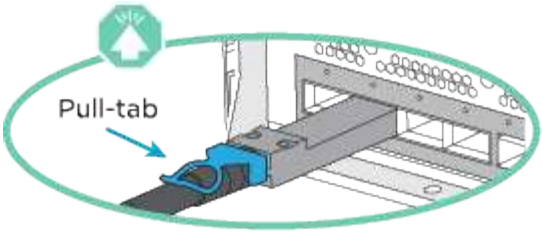

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to cable your controllers to a single shelf:

Animation - Cable the controllers to a single drive shelf

Step	Perform on each controller module
1	Cable controller A to the shelf: 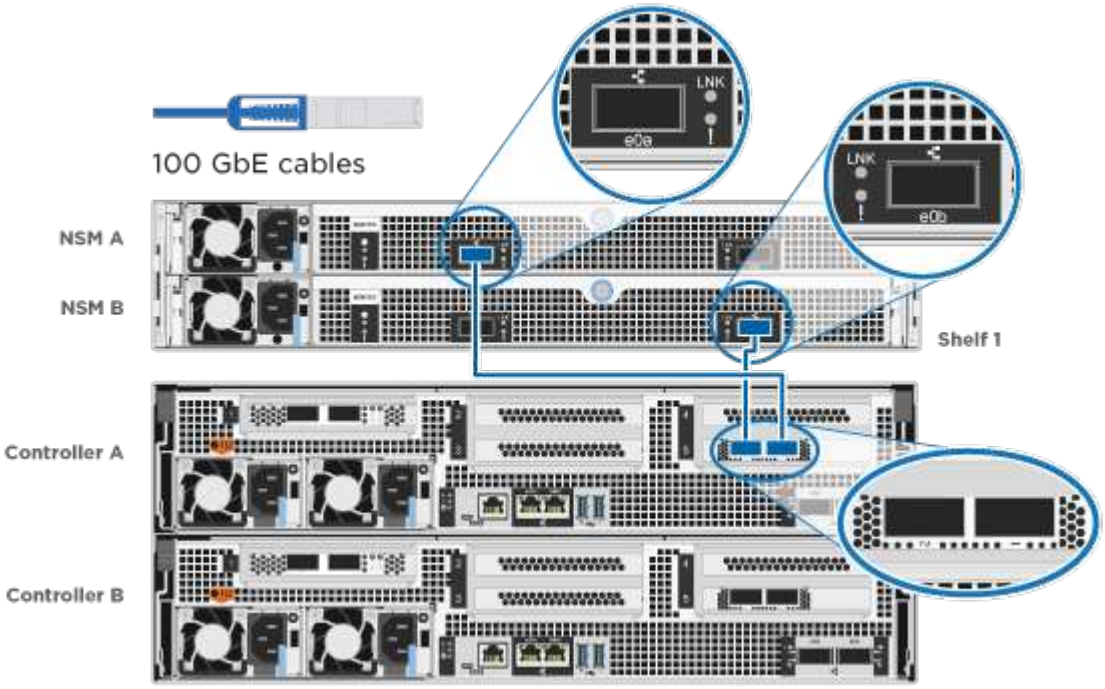

Step	Perform on each controller module
2	Cable controller B to the shelf: 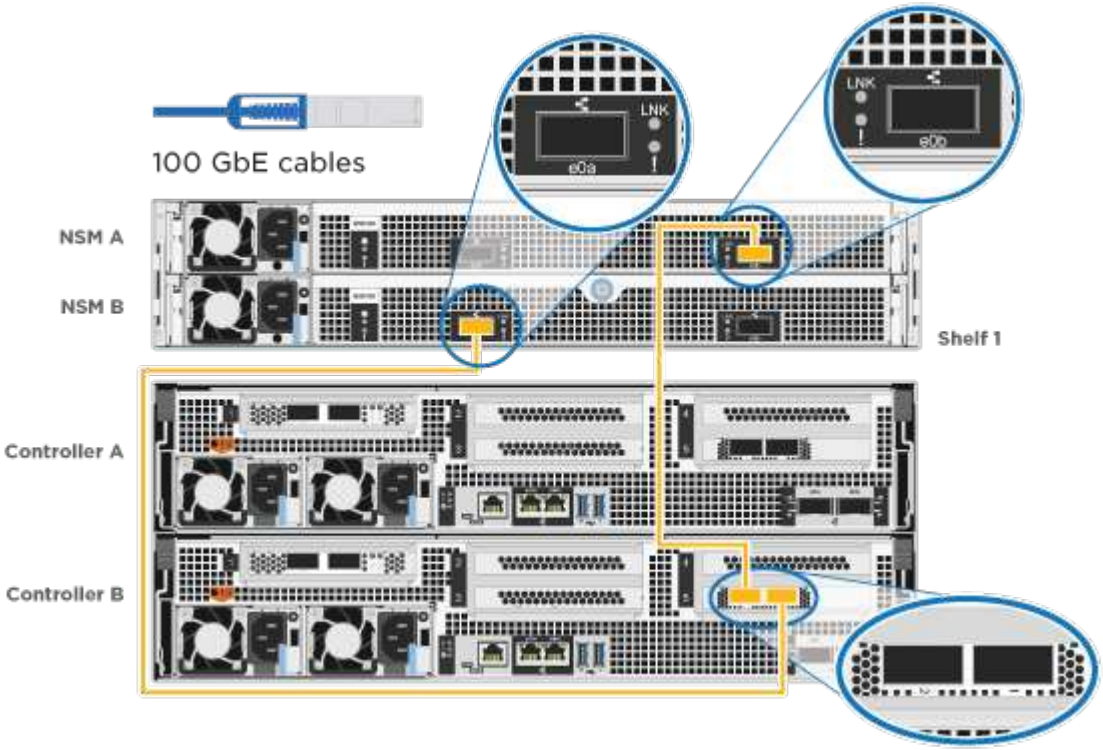

To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 4: Cable the controllers to two drive shelves

You must cable each controller to the NSM modules on both NS224 drive shelves.

Before you begin

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

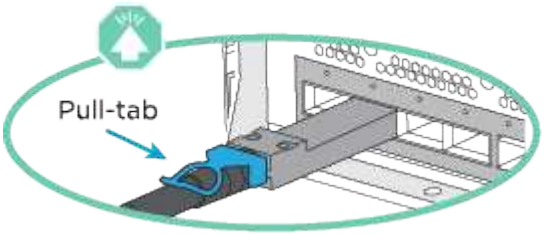

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to cable your controllers to two drive shelves:

Animation - Cable the controllers to two drive shelves

Step	Perform on each controller module
1	Cable controller A to the shelves: 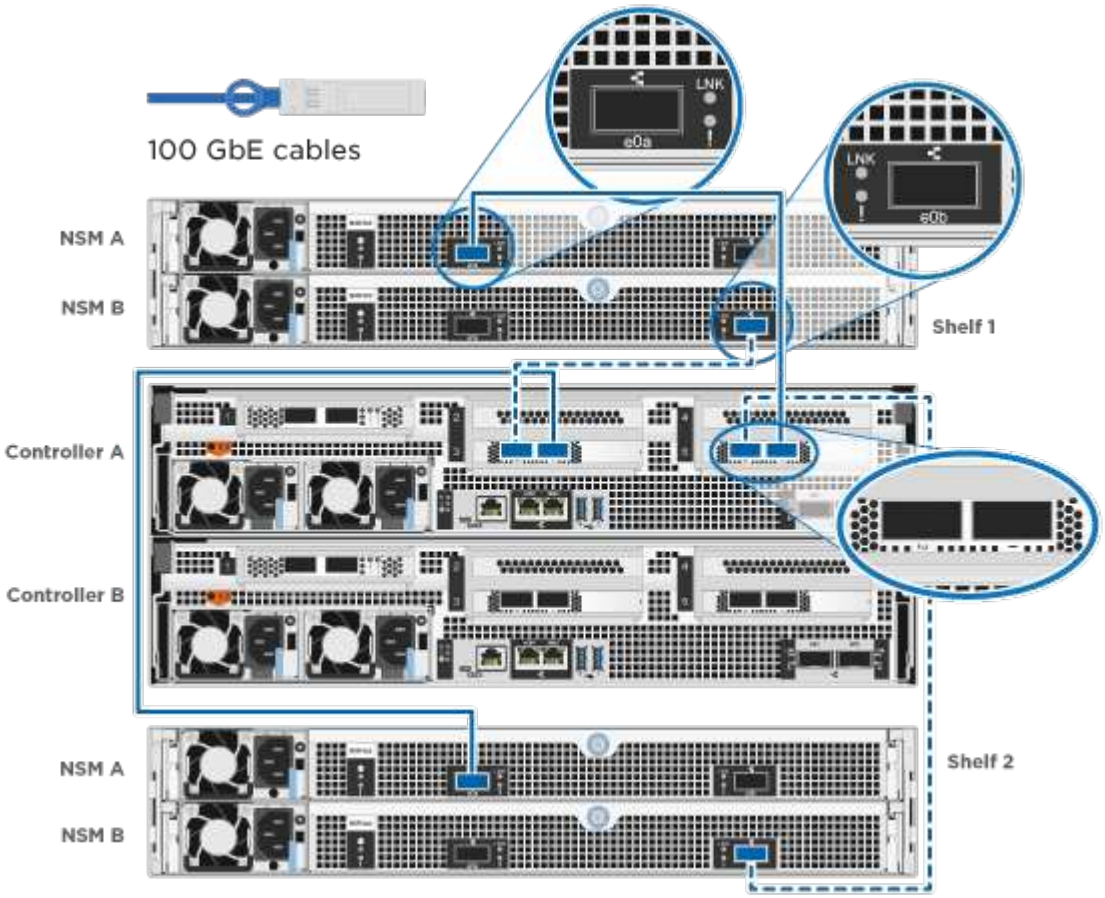

Step	Perform on each controller module
2	Cable controller B to the shelves: 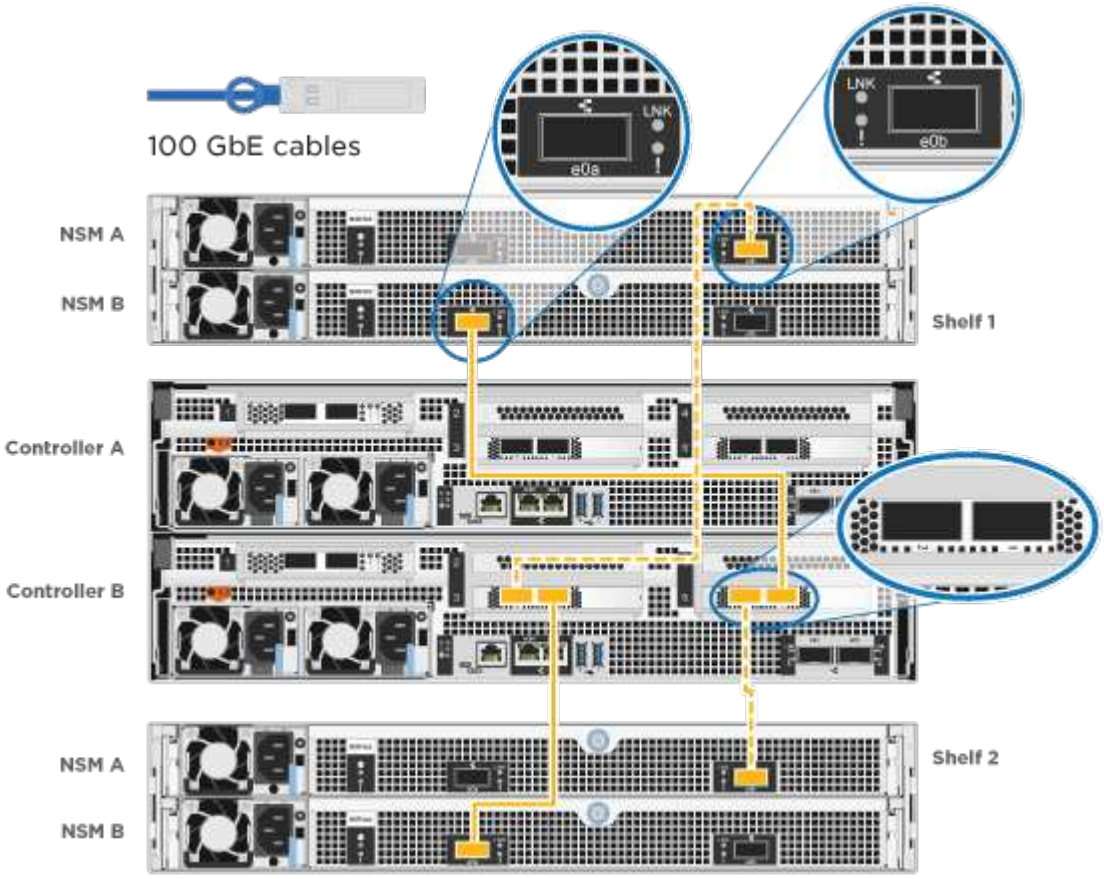

To complete setting up your system, see Step 4: Complete system setup and configuration.

Step 4: Complete system setup and configuration

Complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Complete system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

2. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

3. Use the animation to connect your laptop to the Management switch:

Animation - Connect your laptop to the Management switch

4. Select an ONTAP icon listed to discover:

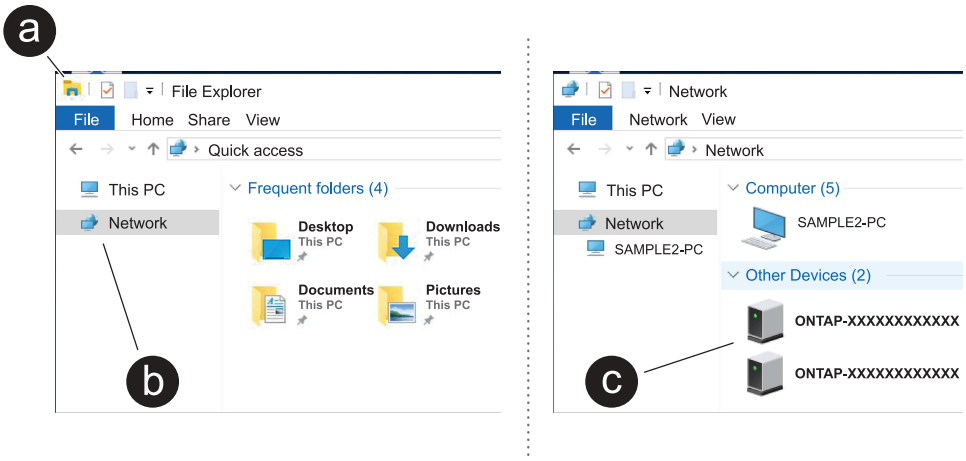

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

      System Manager opens.

5. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
6. Verify the health of your system by running Config Advisor.
7. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Complete system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your system.

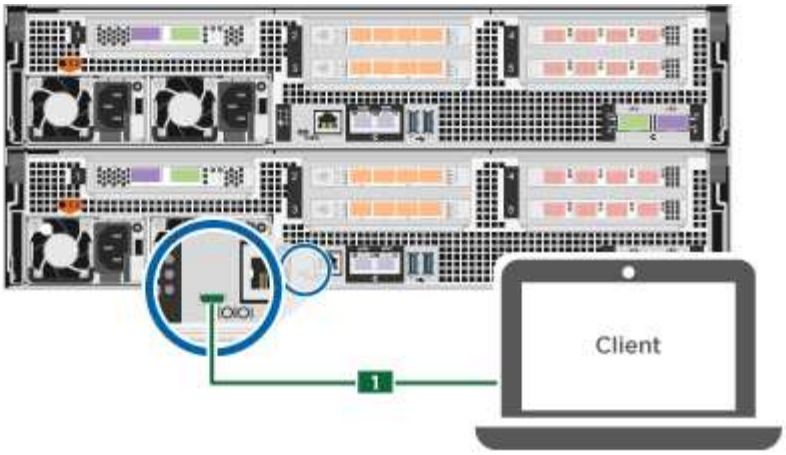

c. Connect the laptop or console to the switch on the management subnet.

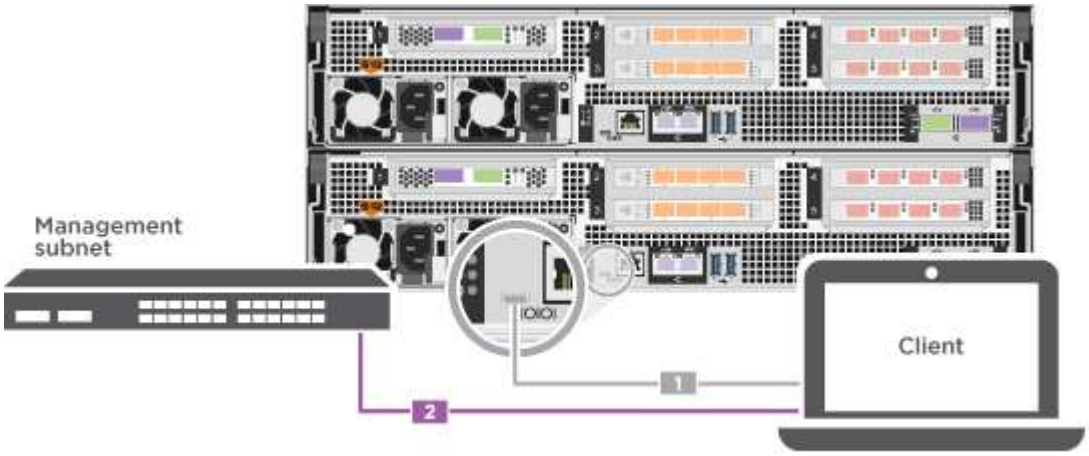

d. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes.

3. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

4. Using System Manager on your laptop or console, configure your cluster:

a. Point your browser to the node management IP address.

> The format for the address is https://x.x.x.x.

b. Configure the system using the data you collected in the ONTAP Configuration Guide.

5. Verify the health of your system by running Config Advisor.
6. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain ASA C800 hardware

Maintain the hardware of your ASA C800 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the ASA C800 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the ASA C800 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Drive	A drive is a device that provides the physical storage media for data.

Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.
NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.
PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - ASA C800

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your ASA C800 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5 **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - ASA C800

Before replacing the boot media in your ASA C800, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - ASA C800

Shut down the impaired controller in your ASA C800 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - ASA C800

The boot media in your ASA C800 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module,

removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

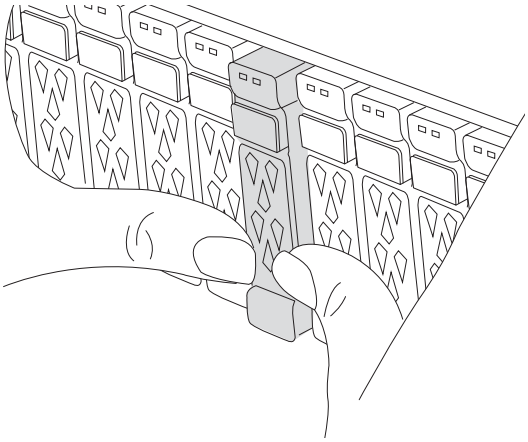

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

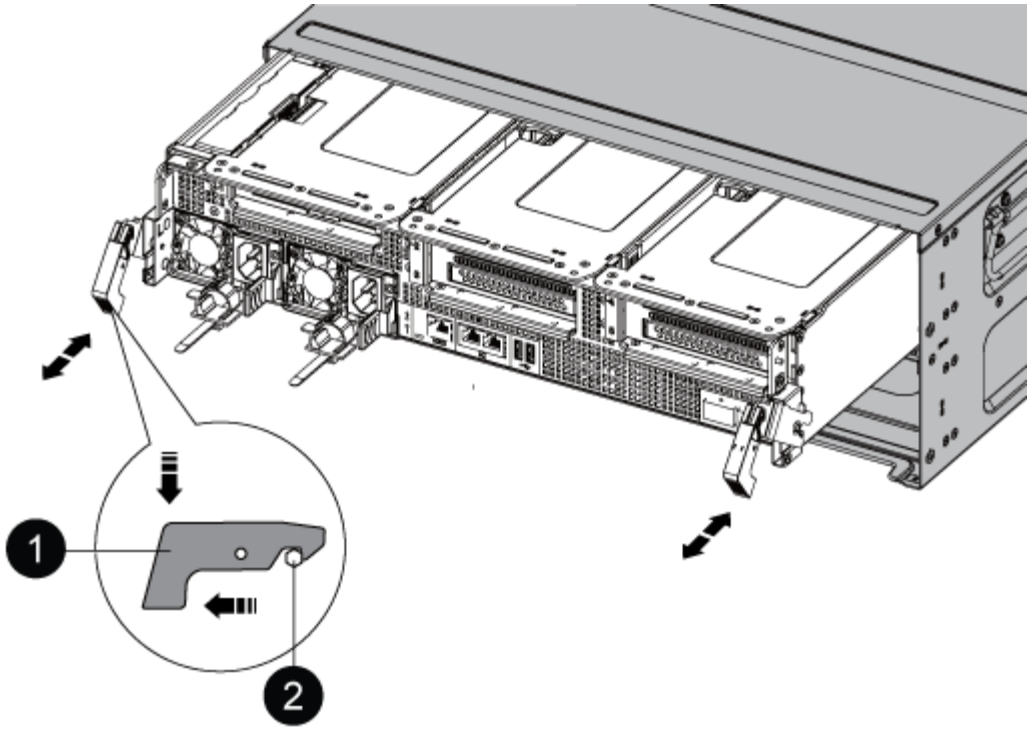

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

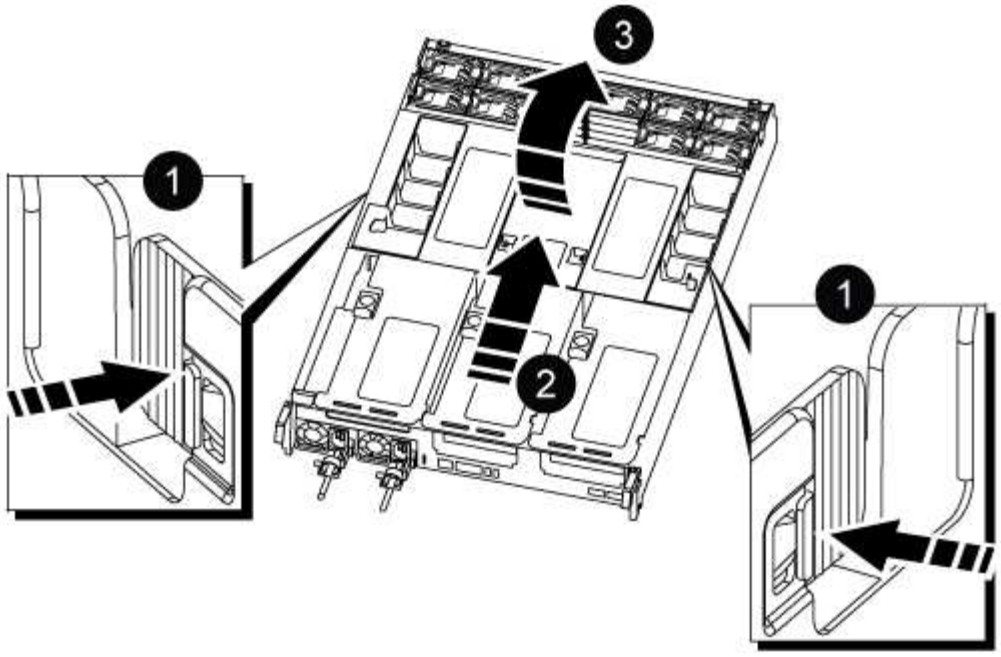

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

10. Locate the boot media in the controller module and replace it:

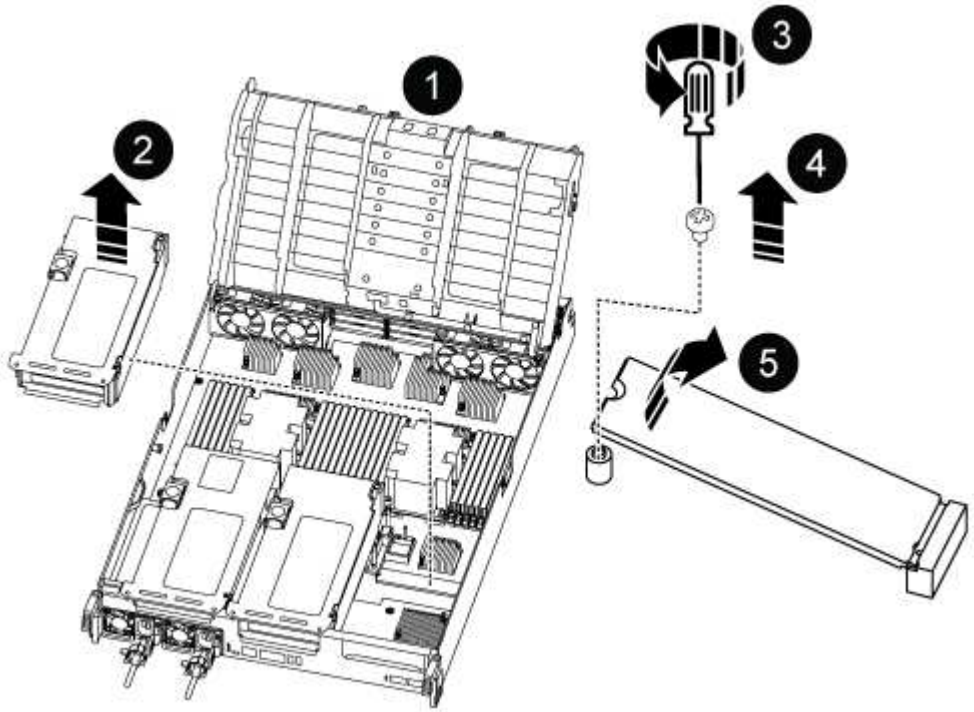

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.

b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.

11. Install the replacement boot media into the controller module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down toward the motherboard.

c. Secure the boot media to the motherboard using the boot media screw.

Do not over-tighten the screw or you might damage the boot media.

12. Reinstall the riser into the controller module.

13. Close the air duct:

a. Rotate the air duct downward.

b. Slide the air duct toward the risers until it clicks into place.

14. Install the controller module:

a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module half-way into the way into the system.

b. Recable the controller module, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

The controller module begins to boot and stops at the LOADER prompt.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - ASA C800

After installing the new boot media device in your ASA C800 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and

determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   `boot_recovery -partner`

   The screen displays the following message:

   `Starting boot media recovery (BMR) process. Press Ctrl-C to abort…`

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

   Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

   Show example of configuration error finding prompts

   ```
   Error when fetching key manager config from partner ${partner_ip}:
   ${status}

   Has key manager been configured on this system

   Is the key manager onboard
   ```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - ASA C800

If a component in your ASAF C800 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - ASA C800

Get started with replacing the boot media in your ASA C800 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - ASA C800

Before replacing the boot media in your ASA C800 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - ASA C800

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - ASA C800

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: System is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media and prepare for manual boot recovery - ASA C800

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

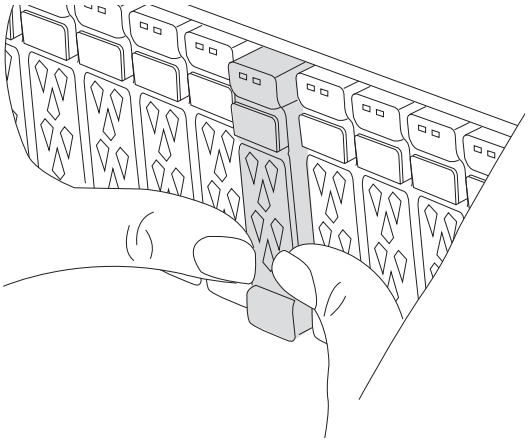

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management

device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

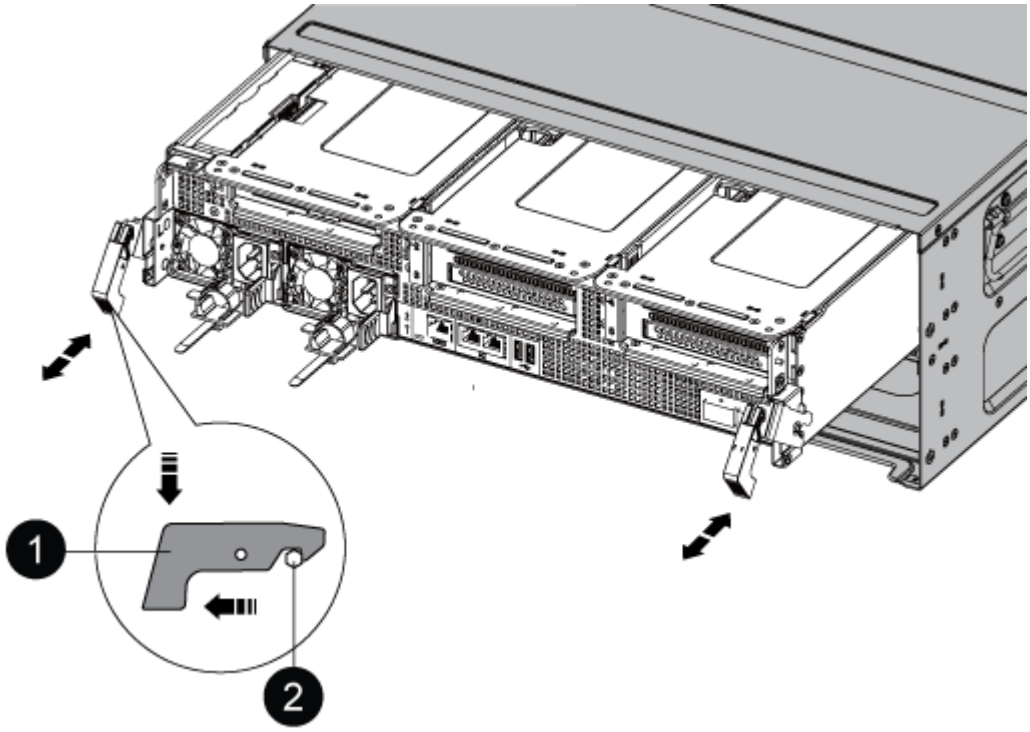

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

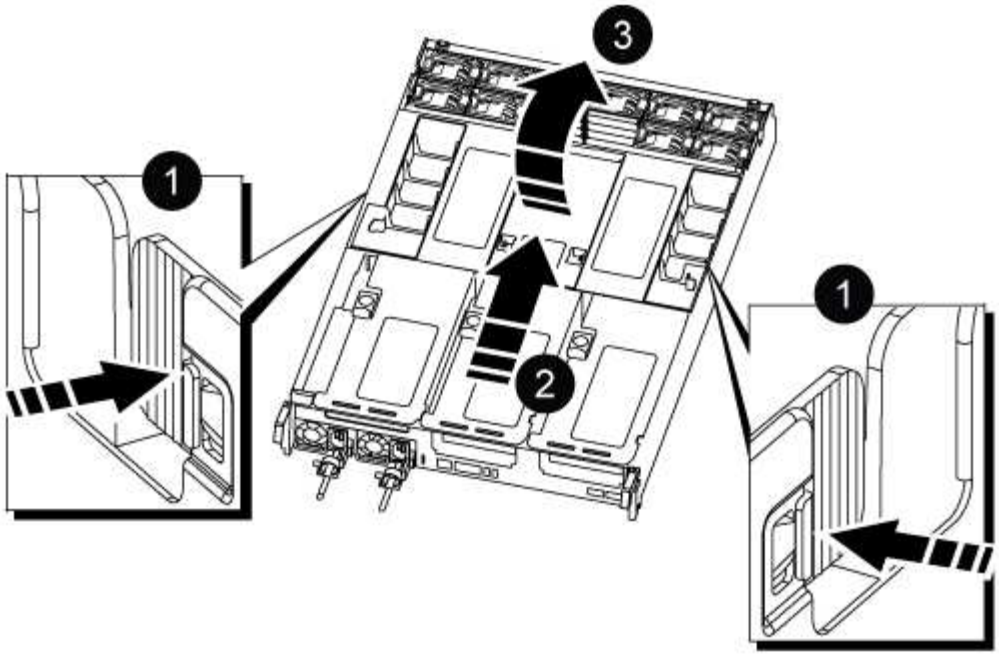

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 2: Replace the boot media

You locate the failed boot media in the controller module by removing Riser 3 on the controller module before you can replace the boot media.

You need a Phillips head screwdriver to remove the screw that holds the boot media in place.

1. Locate the boot media:

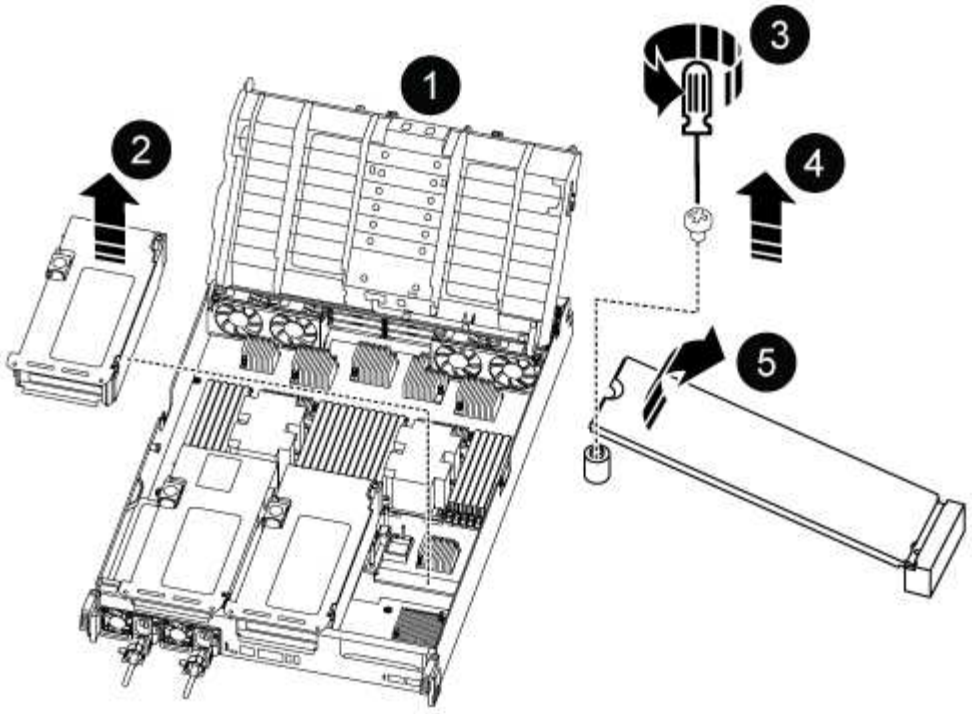

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

2. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.
   b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.
3. Install the replacement boot media into the controller module:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the motherboard.
   c. Secure the boot media to the motherboard using the boot media screw.

      Do not over-tighten the screw or you might damage the boot media.
4. Reinstall the riser into the controller module.

5. Close the air duct:
    a. Rotate the air duct downward.
    b. Slide the air duct toward the risers until it clicks into place.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed is without a boot image so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
    - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
    - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
    a. Download the service image to your work space on your laptop.
    b. Unzip the service image.

        If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

        There are two folders in the unzipped service image file:

        - boot
        - efi

    c. Copy the efi folder to the top directory on the USB flash drive.

        If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

    The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

    a. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct:
    a. Swing the air duct all the way down to the controller module.
    b. Slide the air duct toward the risers until the locking tabs click into place.

c. Inspect the air duct to make sure that it is properly seated and locked into place.

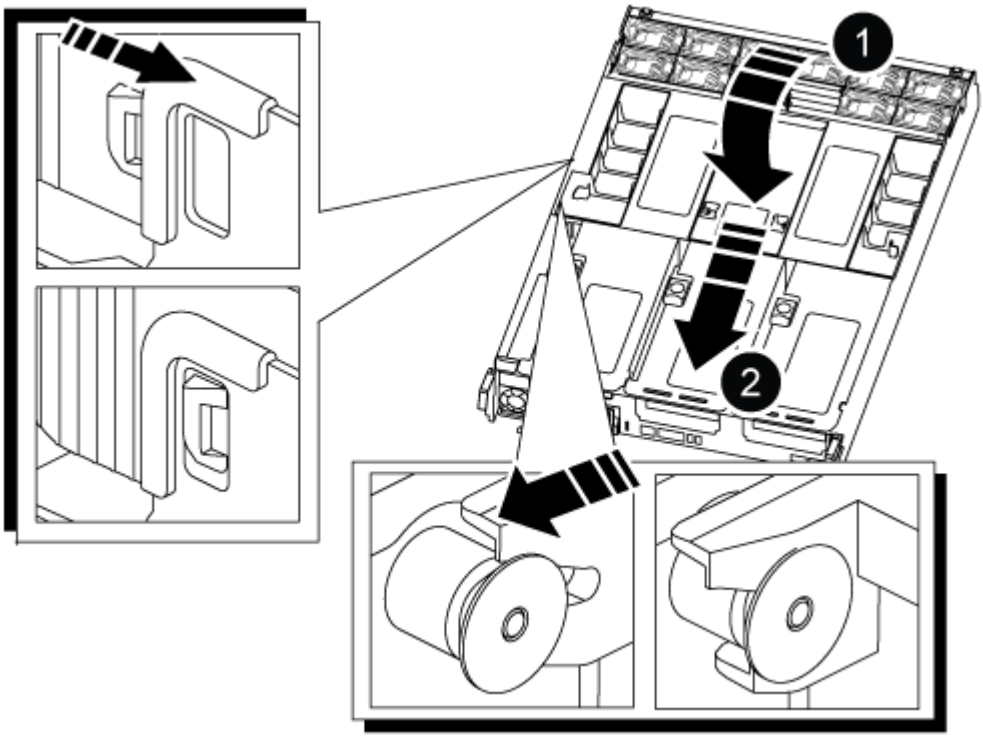

1	Air duct
2	Risers

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.

5. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

6. Gently push the controller module all the way into the system until the controller module locking hooks begin to rise, firmly push on the locking hooks to finish seating the controller module, and then swing the locking hooks into the locked position over the pins on the controller module.
7. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the

controller to boot to LOADER.

Manual boot media recovery from a USB drive - ASA C800

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore encryption - ASA C800

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the manual boot recovery procedure

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - ASA C800

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Replace the chassis - ASA C800

To replace the chassis, you must move the bezel, controller modules, and NVMe drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in its place.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - ASA C800

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - ASA C800

Move the power supplies, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

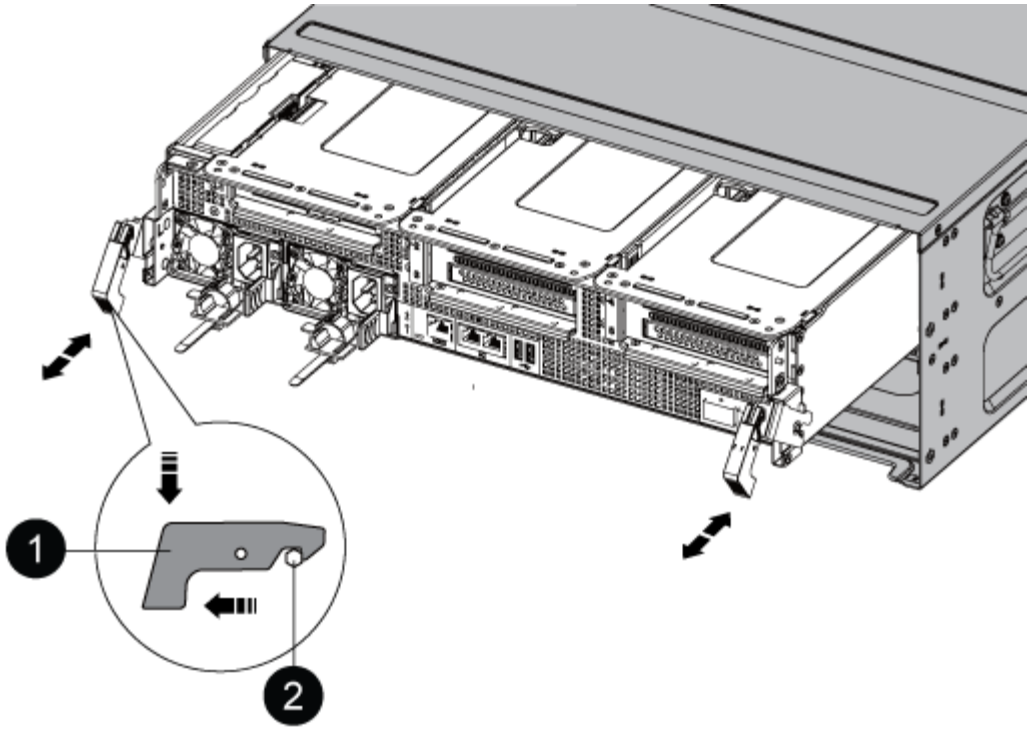

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.

3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.

5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process by pressing `Ctrl-C`.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - ASA C800

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - ASA C800

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Do not downgrade the BIOS version of the *replacement* controller to match the partner controller or the old controller module.

Shut down the impaired controller - ASA C800

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   `cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - ASA C800

To replace the controller, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

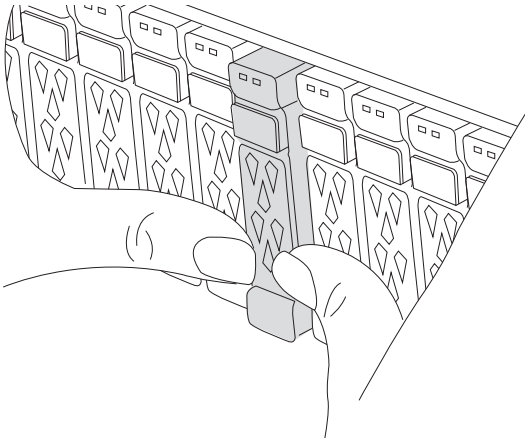

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

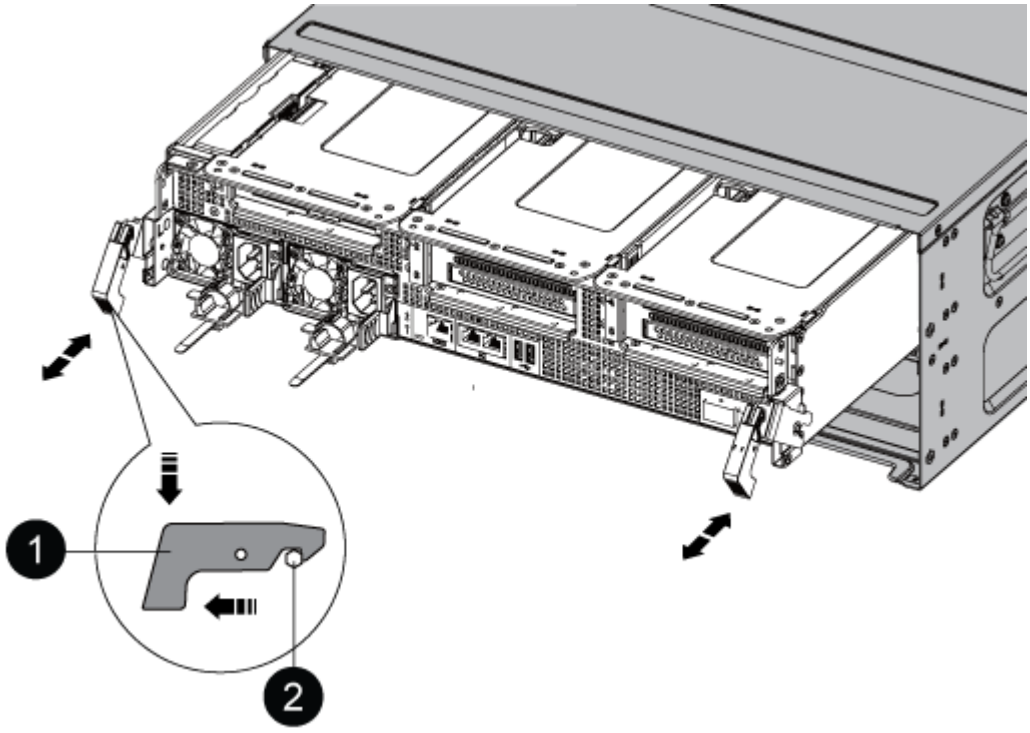

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
   .

9. Place the controller module on a stable, flat surface.
10. Open the controller module air duct:
    a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
    b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

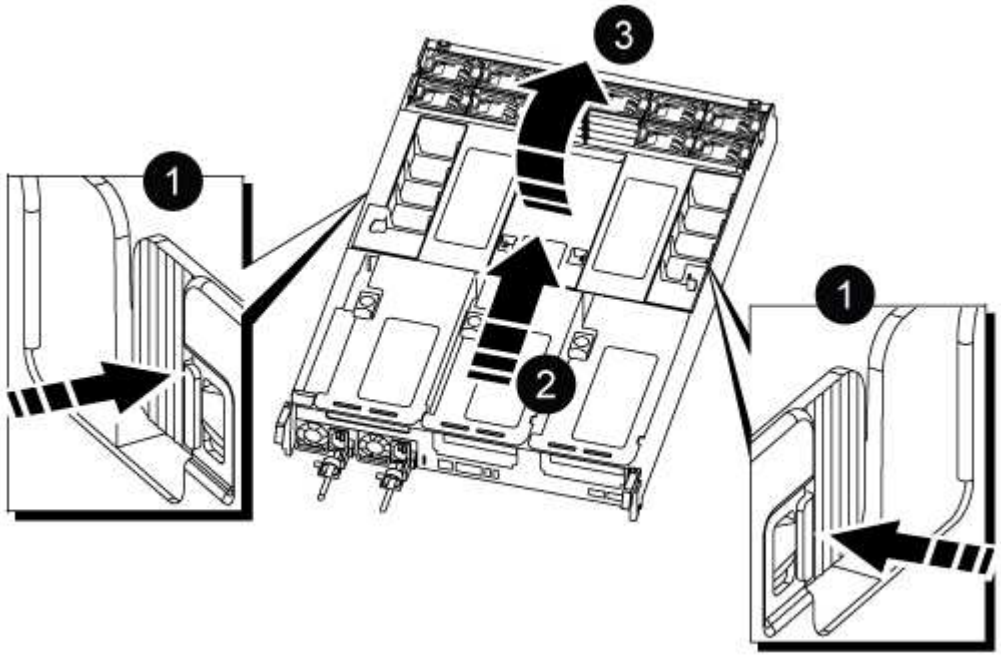

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 2: Move the power supplies

You must move the power supplies from the impaired controller module to the replacement controller module when you replace a controller module.

1. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

> The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

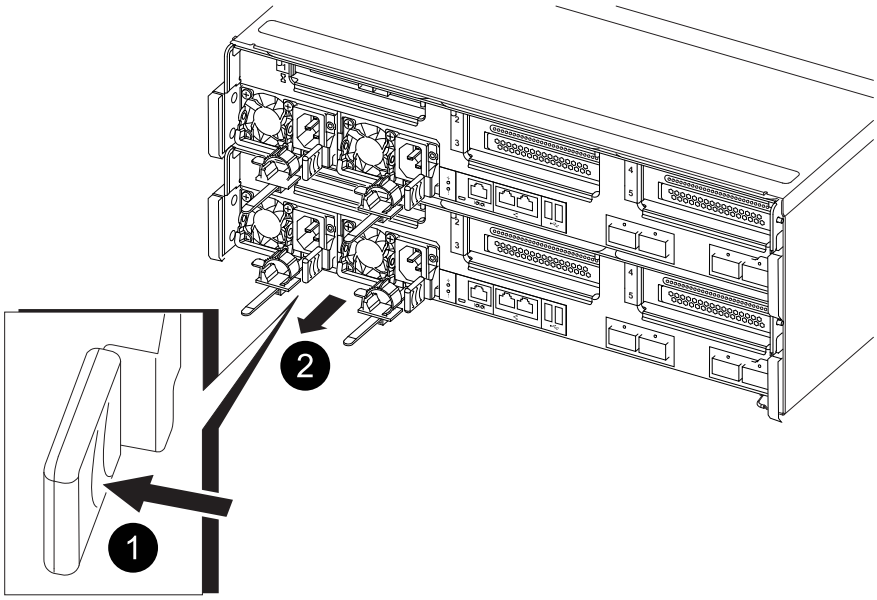

1	Blue power supply locking tab
2	Power supply

2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

> To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

You must move the fans from the impaired controller module to the replacement module when replacing a failed controller module.

1. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

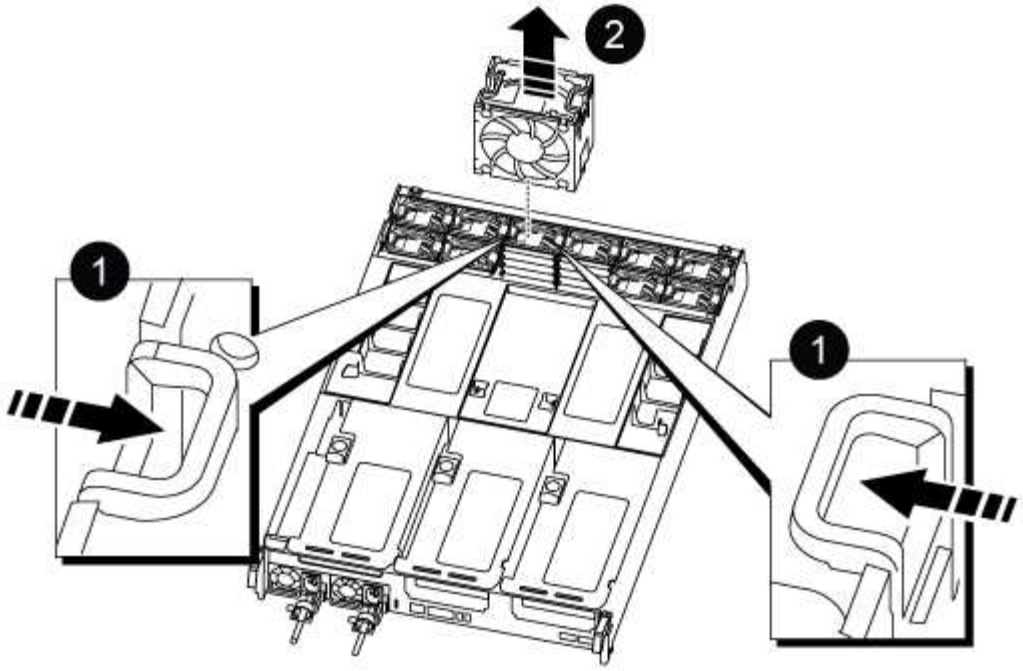

1	Fan locking tabs
2	Fan module

2. Move the fan module to the replacement controller module, and then install the fan module by aligning its edges with the opening in the controller module, and then sliding the fan module into the controller module until the locking latches click into place.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the NVDIMM battery

When replacing the controller module, you must move the NVRAM battery from the impaired controller module to the replacement controller module

1. Open the air duct cover and locate the NVDIMM battery in the riser.

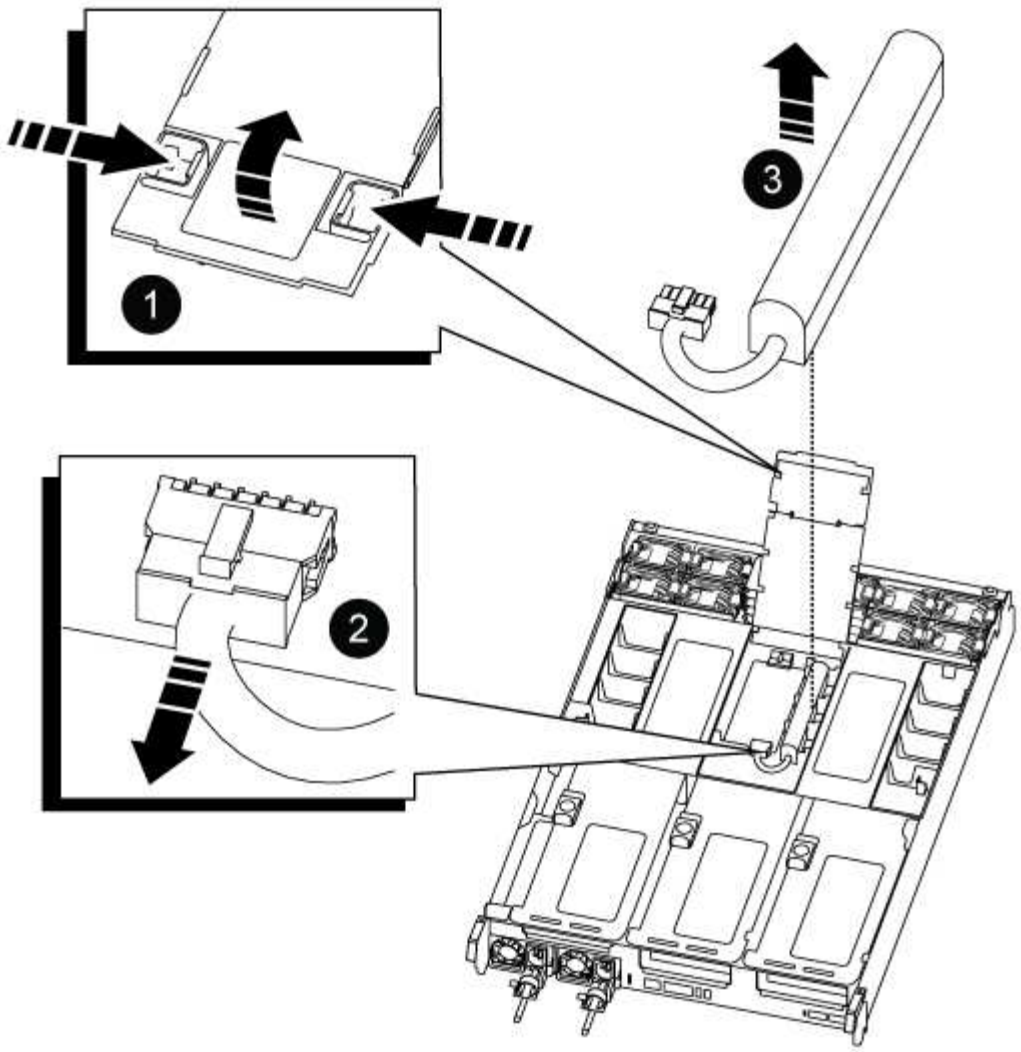

1	Air duct riser
2	NVDIMM battery plug
3	NVDIMM battery pack

Attention: The NVDIMM battery control board LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
3. Grasp the battery and lift the battery out of the air duct and controller module.
4. Move the battery pack to the replacement controller module and then install it in the NVDIMM air duct:
   a. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   b. Plug the battery plug into the riser socket and make sure that the plug locks into place.

Step 5: Remove the PCIe risers

As part of the controller replacement process, you must remove the PCIe modules from the impaired controller module. You must install them into the same location in the replacement controller module once the NVDIMMS and DIMMs have moved to the replacement controller module.

1. Remove the PCIe riser from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

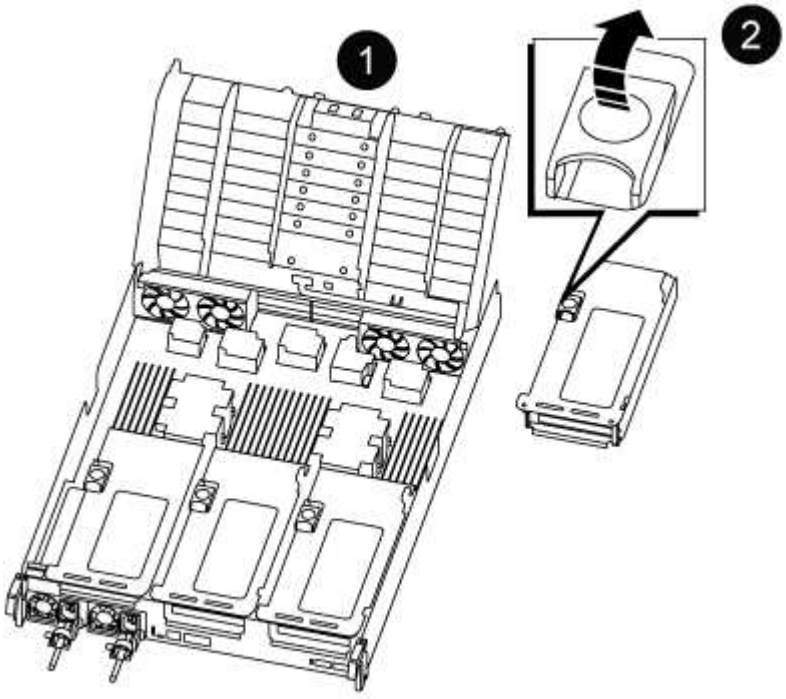

1	Air duct
2	Riser 1 (left riser), Riser 2 (middle riser), and 3 (right riser) locking latches

2. Repeat the preceding step for the remaining risers in the impaired controller module.
3. Repeat the above steps with the empty risers in the replacement controller and put them away.

Step 6: Move system DIMMs

To move the DIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

1. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.

2. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

3. Locate the slot where you are installing the DIMM.
4. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

5. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
6. Repeat these steps for the remaining DIMMs.

Step 7: Move the NVDIMMs

To move the NVDIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

1. Locate the NVDIMMs on your controller module.

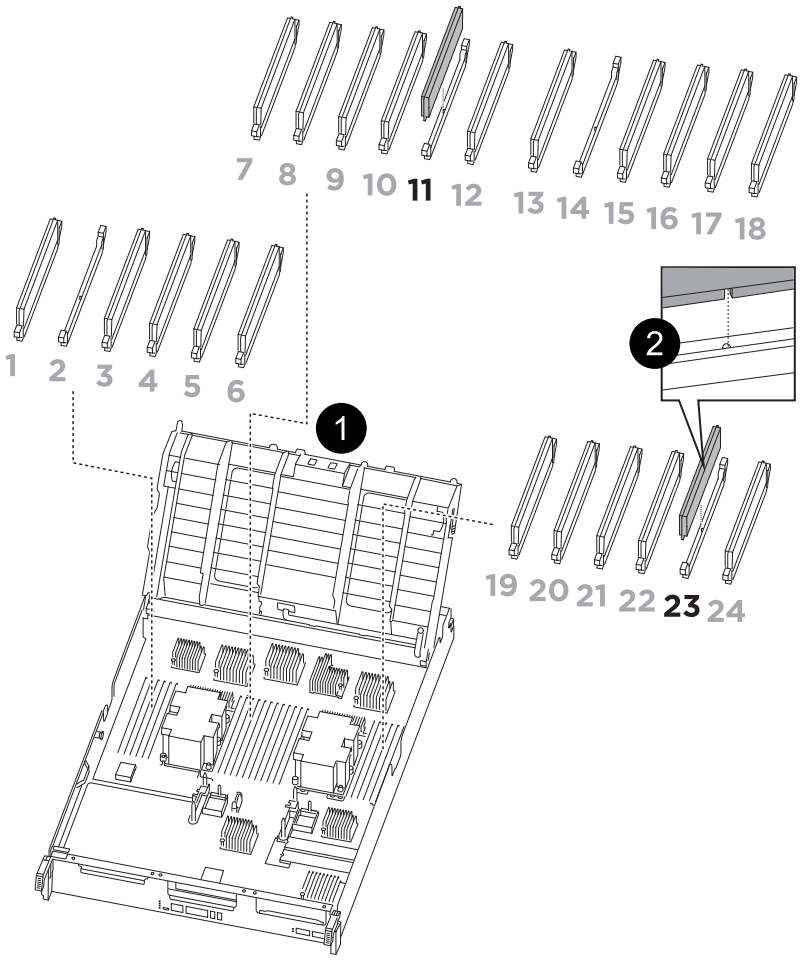

- NVDIMM: SLOTS 11 & 23

1	Air duct
2	NVDIMMs

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

4. Locate the slot where you are installing the NVDIMM.
5. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
7. Repeat the preceding steps to move the other NVDIMM.

Step 8: Move the boot media

You must move the boot media device from the impaired controller and install it in the replacement controller.

The boot media is located under Riser 3.

1. Locate the boot media:

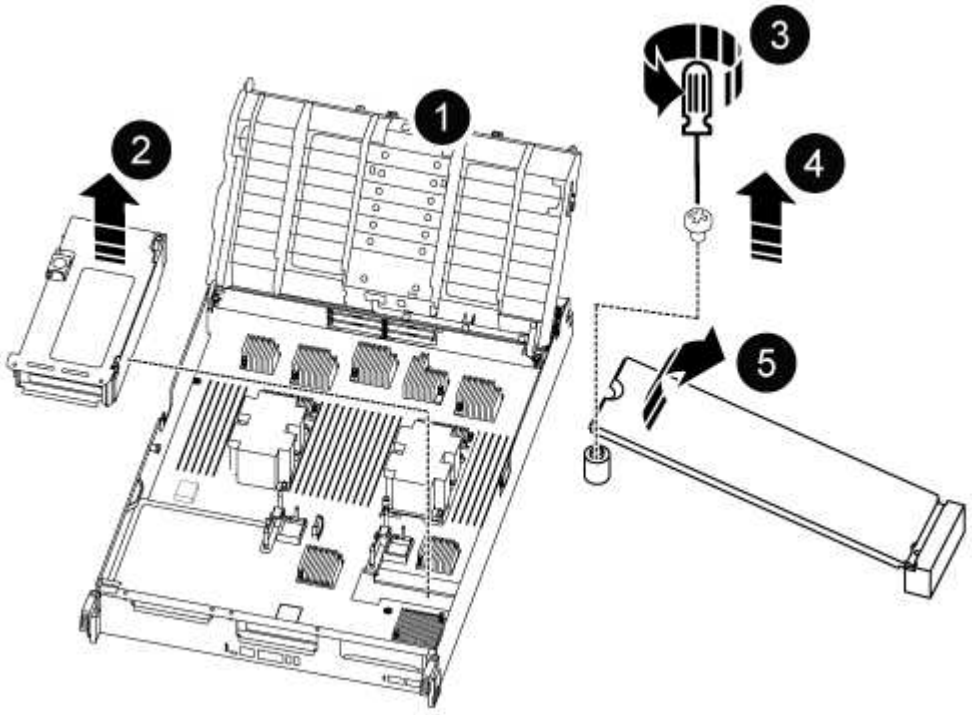

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

2. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.

b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.

3. Move the boot media to the new controller module and install it:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the motherboard.
   c. Secure the boot media to the motherboard using the boot media screw.

      Do not over-tighten the screw or you might damage the boot media.

Step 9: Install the PCIe risers

You install the PCIe risers in the replacement controller module after moving the DIMMs, NVDIMMs, and boot media.

1. Install the riser into the replacement controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

   d. Reinsert any SFP or QSFP modules that were removed from the PCIe cards.
2. Repeat the preceding step for the remaining PCIe risers.

Step 10: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis and then boot it to Maintenance mode.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

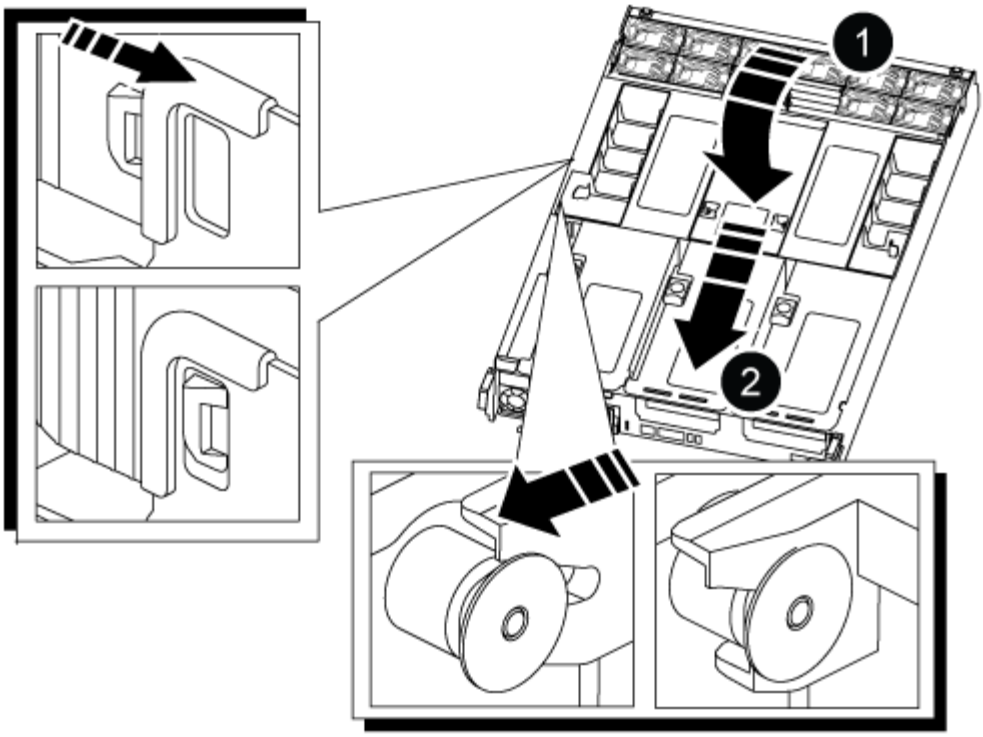

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:

   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

5. Plug the system cables and transceiver modules into the controller module and reinstall the cable management device.

6. Plug the power cables into the power supplies and reinstall the power cable retainers.

   The controller module begins to boot as soon as it is connected to power. Be prepared to interrupt the boot process.

   If your system has DC power supplies, make sure the thumbscrews on the power supply cable are tight.

Restore and verify the system configuration - ASA C800

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the chassis

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA`

state: `ha-config show`

The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - ASA C800

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and

then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - ASA C800

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - ASA C800

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

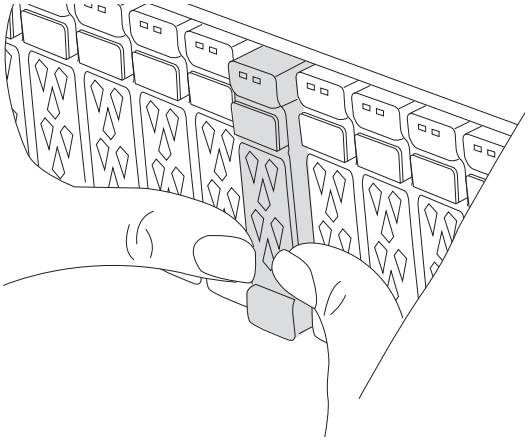

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

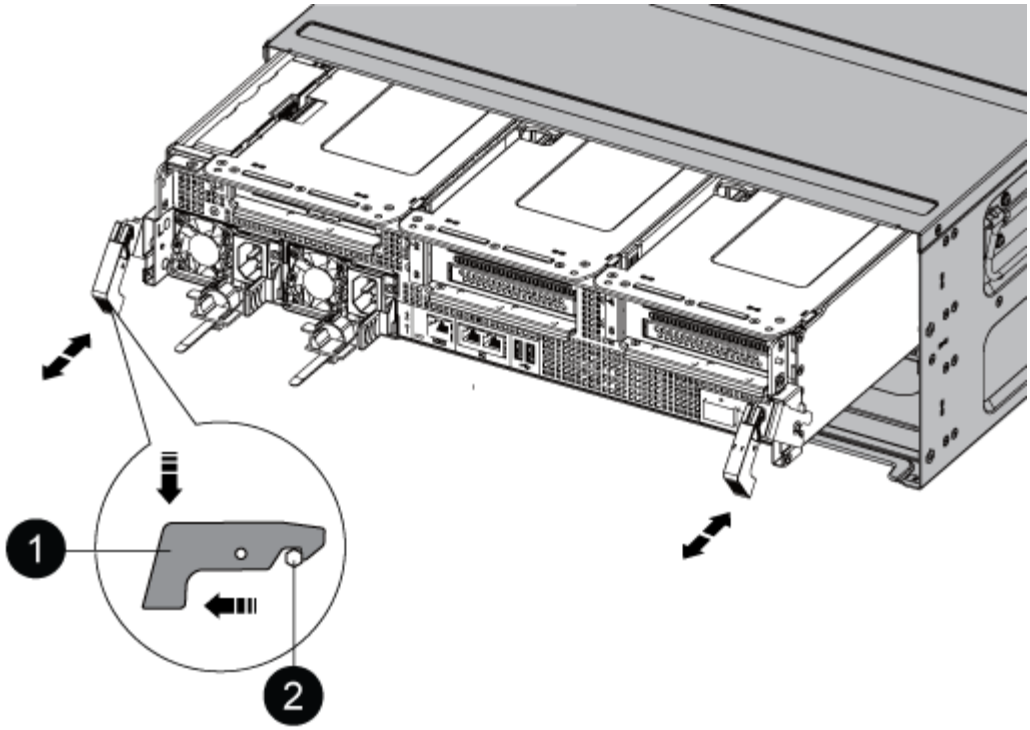

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

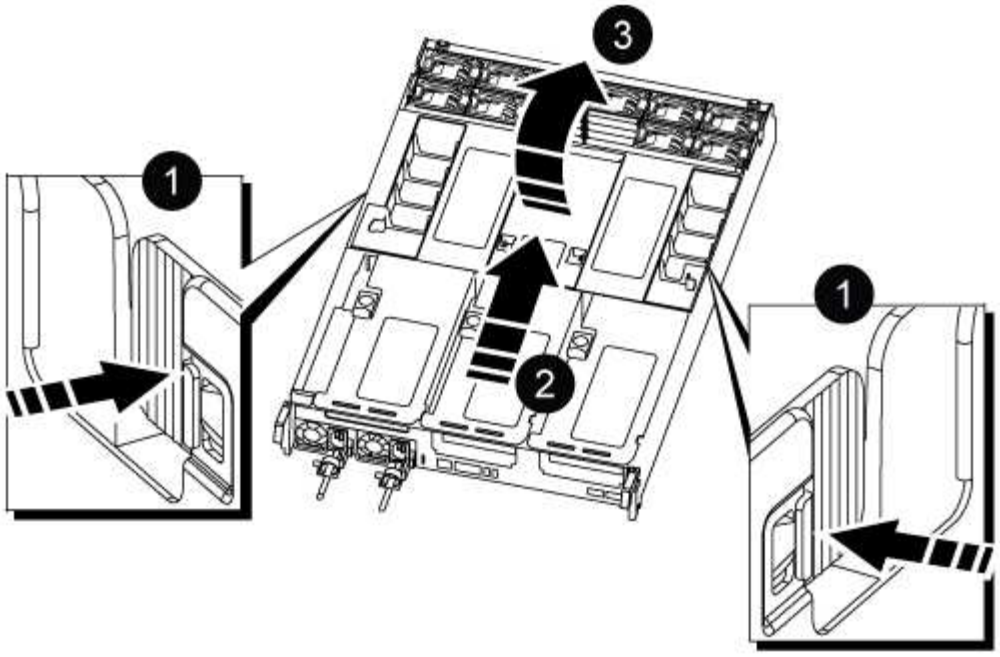

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace a DIMM

To replace a DIMM, you must locate it in the controller module using the DIMM map label on top of the air duct and then replace it following the specific sequence of steps.

1. When removing a DIMM, unlock the locking latch on the applicable riser, and then remove the riser.

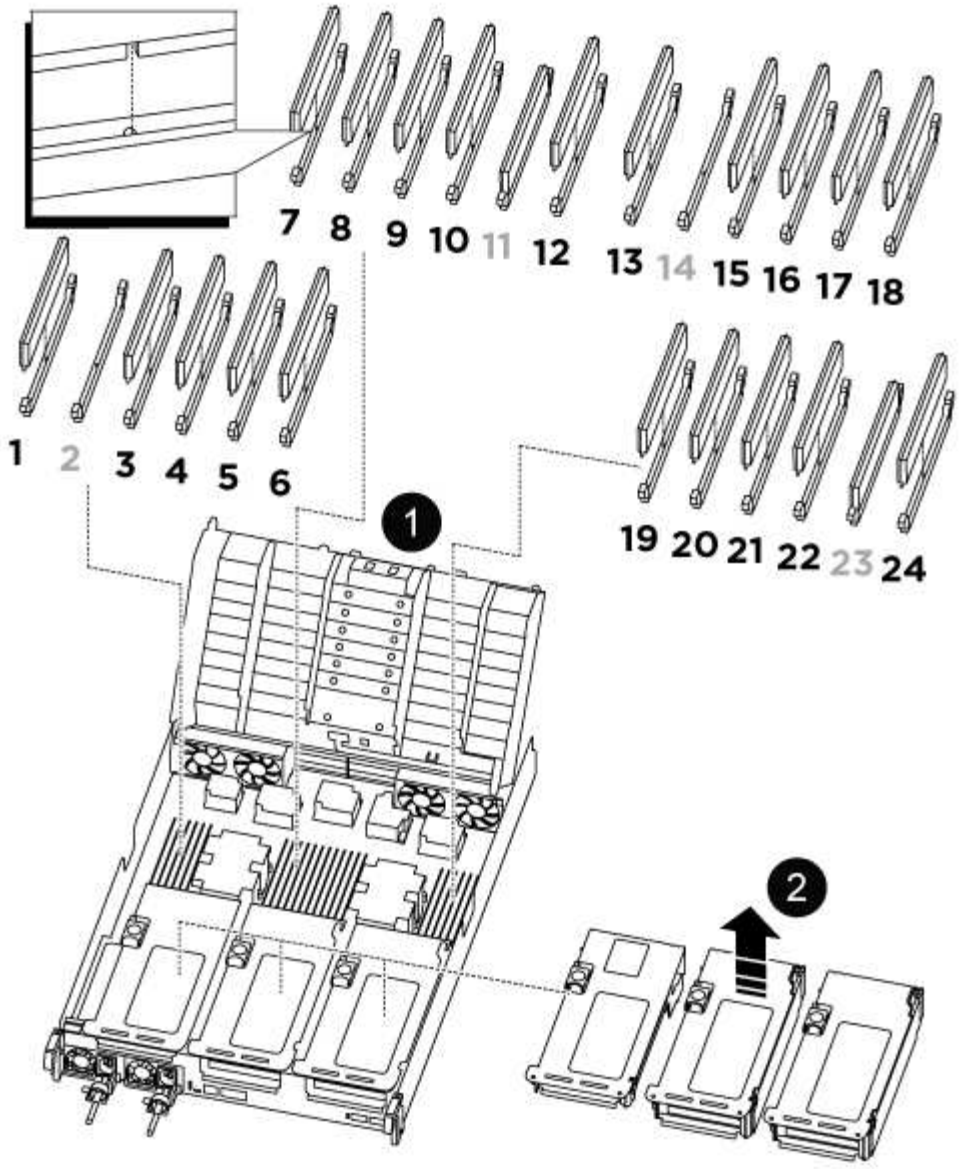

1	Air duct cover
2	Riser 1 and DIMM bank 1, and 3-6
Riser 2 and DIMM bank 7-10, 12-13, and 15-18	Riser 3 and DIMM 19 -22 and 24

Note: Slot 2 and 14 are left empty. Do not attempt to install DIMMs into these slots.

2. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

4. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

5. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Reinstall any risers that you removed from the controller module.
8. Close the air duct.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

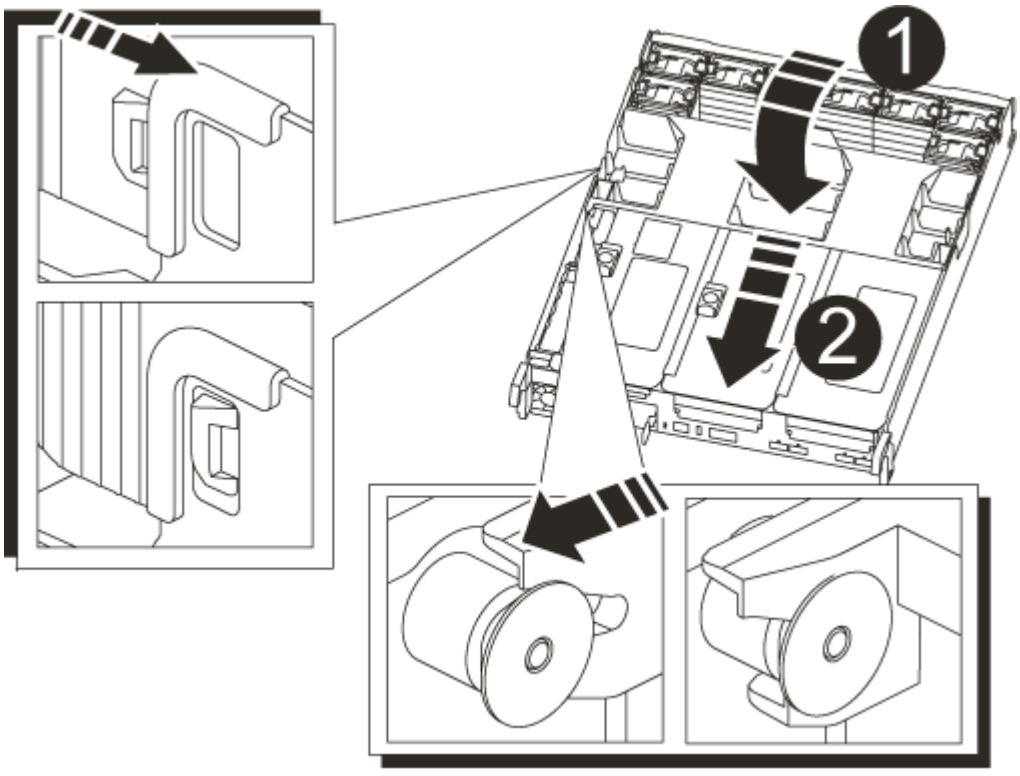

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - ASA C800

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before

replacing a drive.

- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  > Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan - ASA C800

To replace a fan, remove the failed fan module and replace it with a new fan module.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a fan module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

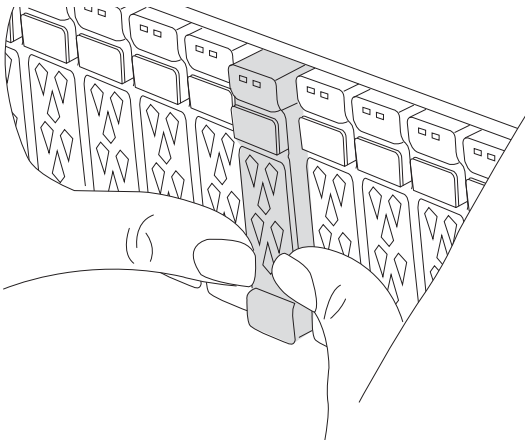

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

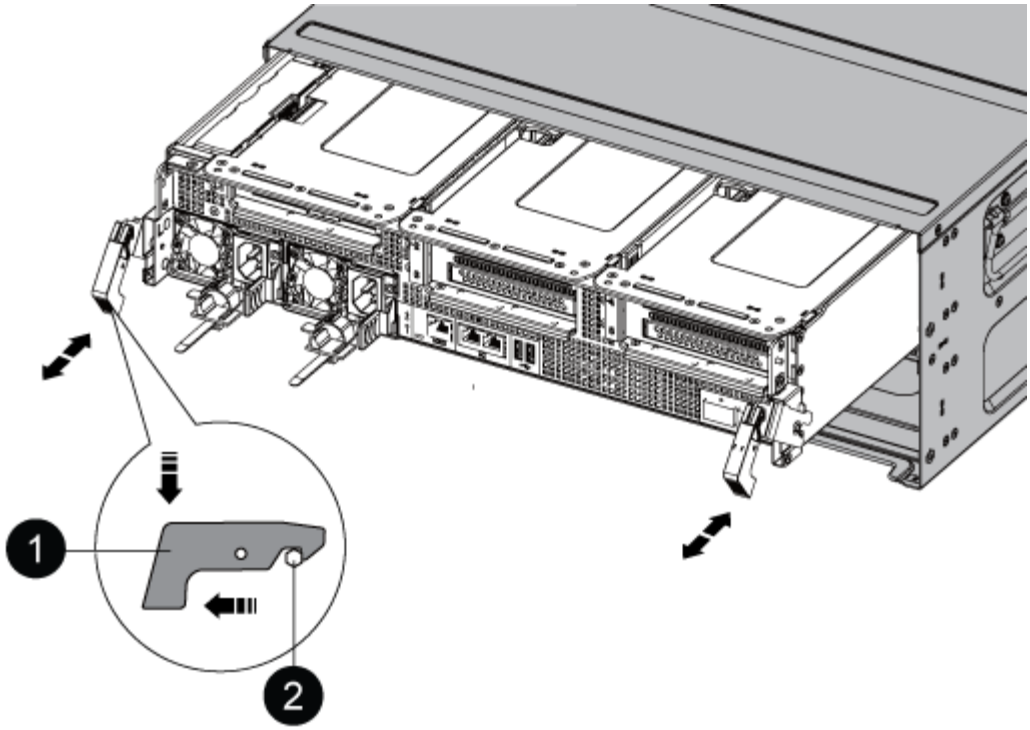

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Set the controller module aside in a safe place.

Step 3: Replace a fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
2. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

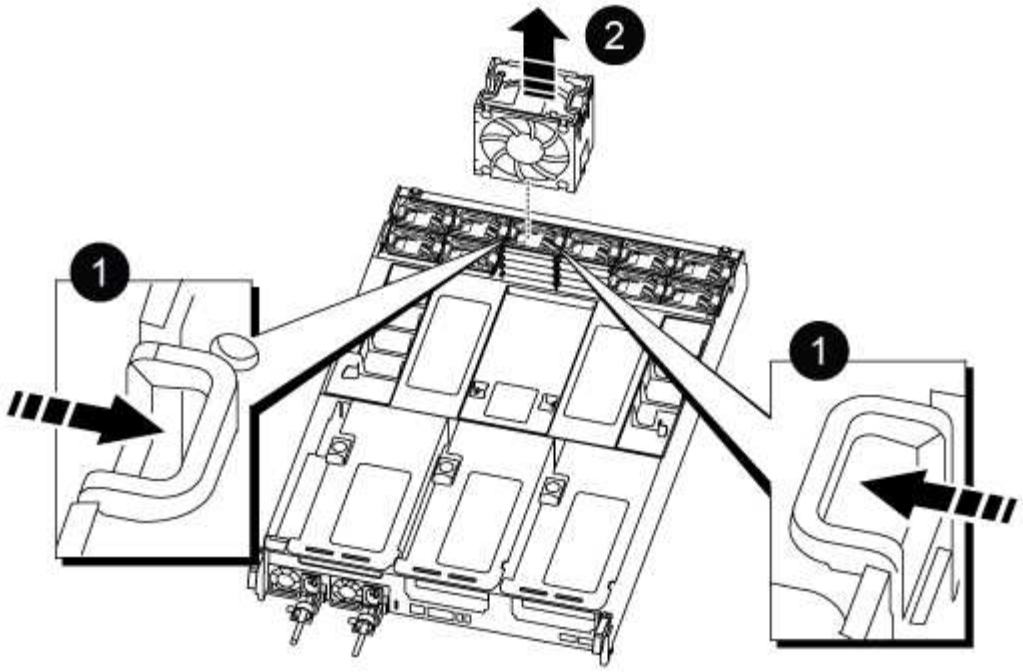

1	Fan locking tabs
2	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module until the locking latches click into place.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.
3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -controller local -auto-giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an NVDIMM - ASA C800

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

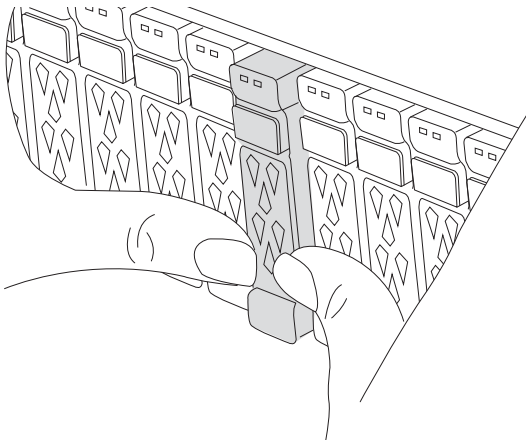

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.

5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

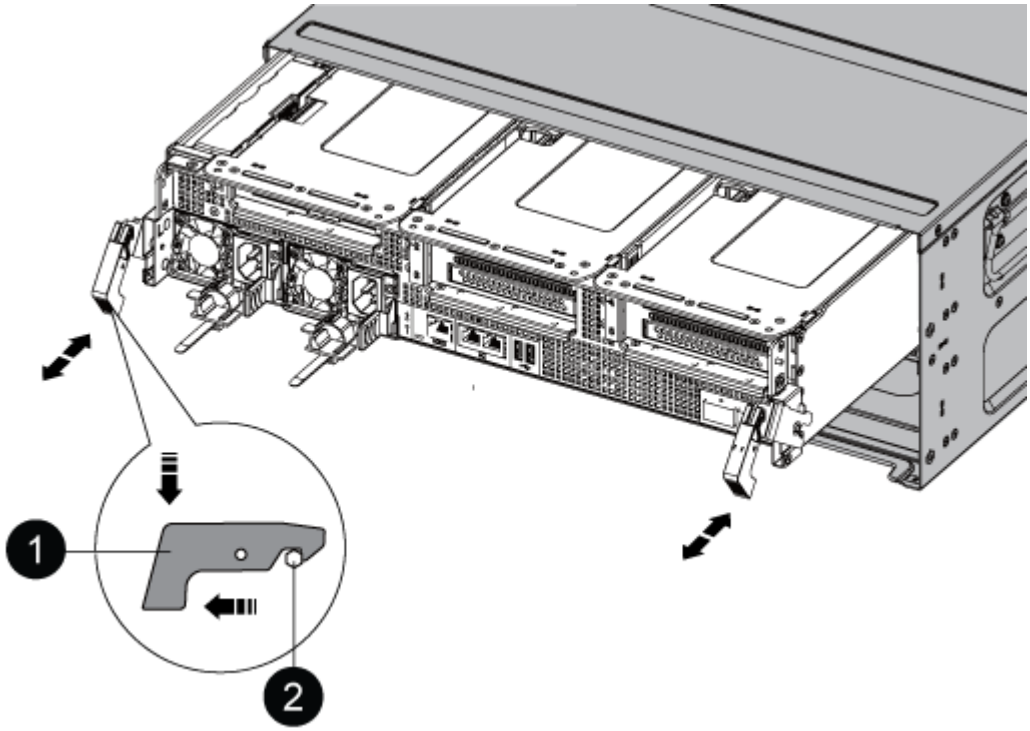

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:

   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

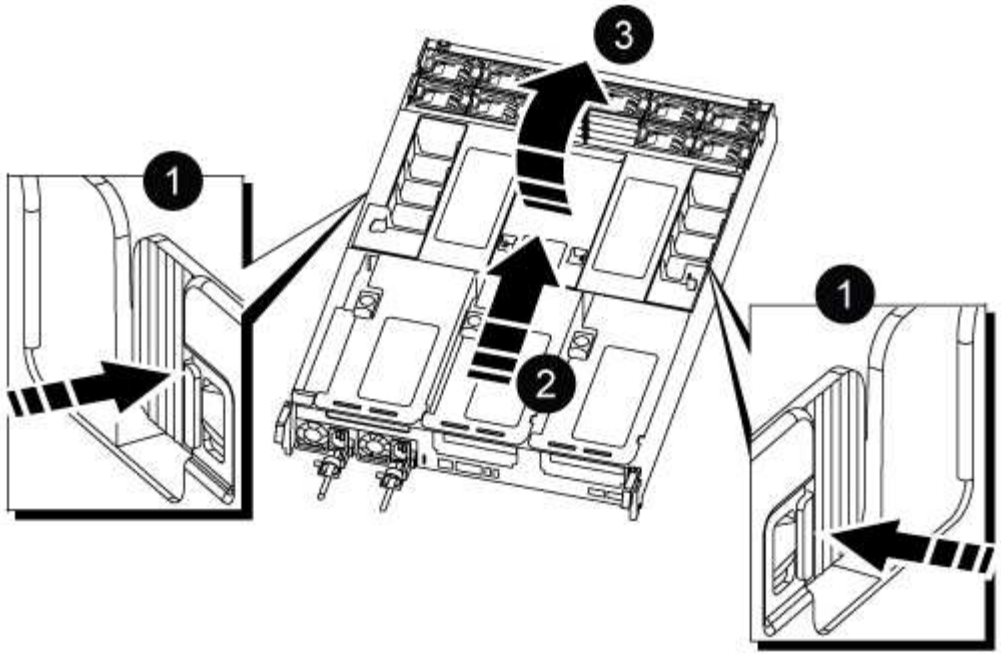

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the NVDIMM map label on top of the air duct, and then replace it following the specific sequence of steps.

1. If you are removing or moving an NVDIMM, unlock the locking latch on the riser, and then remove the applicable riser.

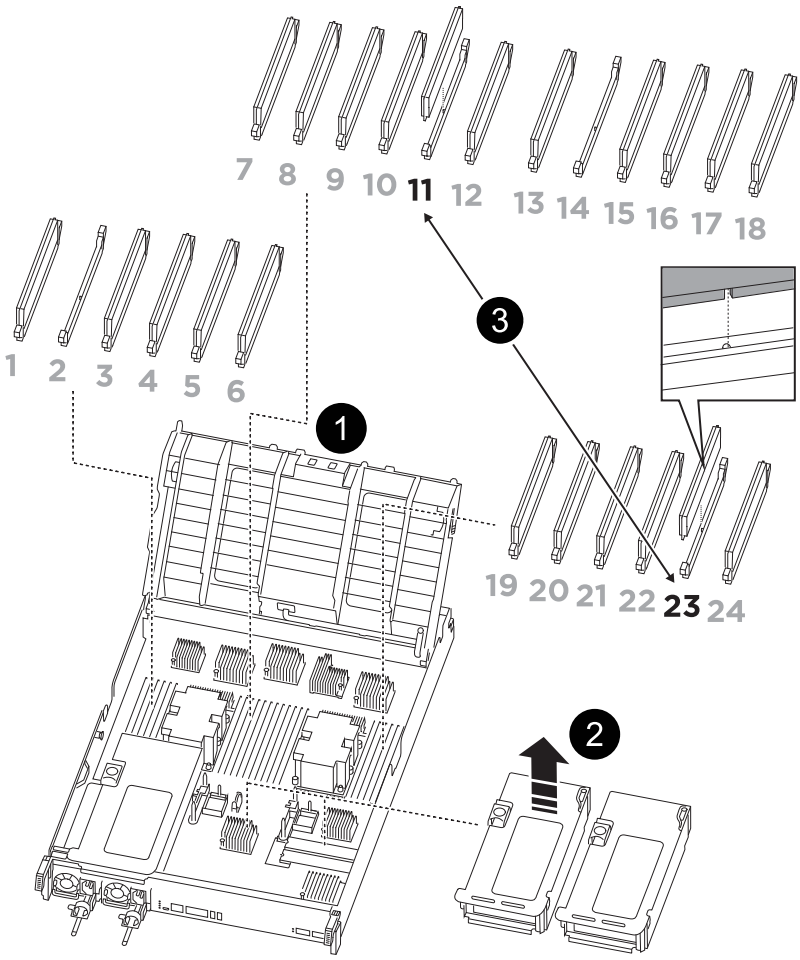

1	Air duct cover
2	Riser 2
3	NVDIMM in slots 11 and 23

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

4. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

5. Locate the slot where you are installing the NVDIMM.

6. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
8. Reinstall any risers that you removed from the controller module.
9. Close the air duct.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

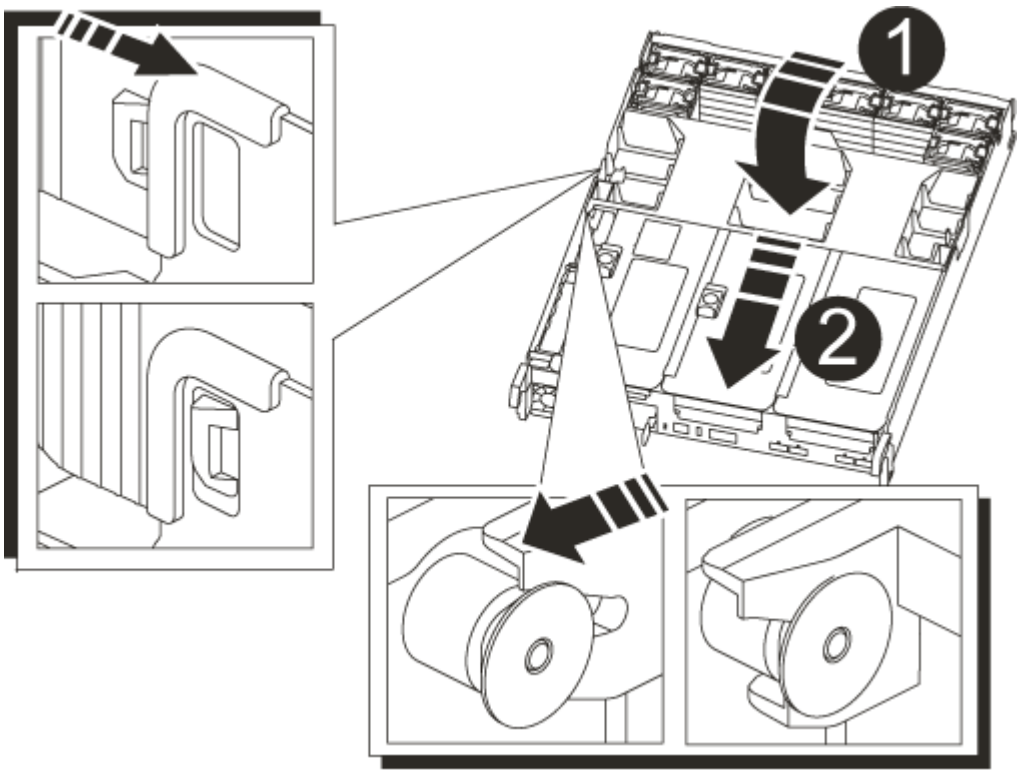

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - ASA C800

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be

resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

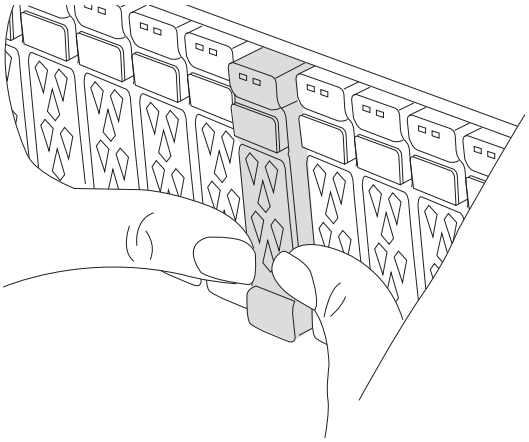

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

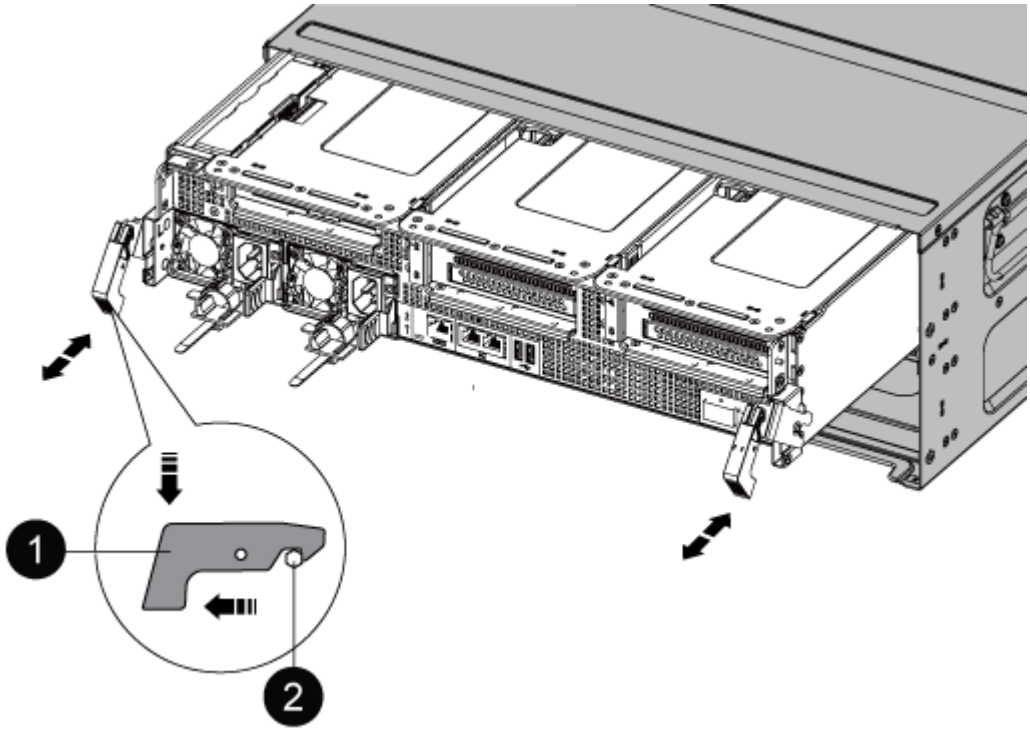

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Set the controller module aside in a safe place.

Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module.

1. Open the air duct cover and locate the NVDIMM battery in the riser.

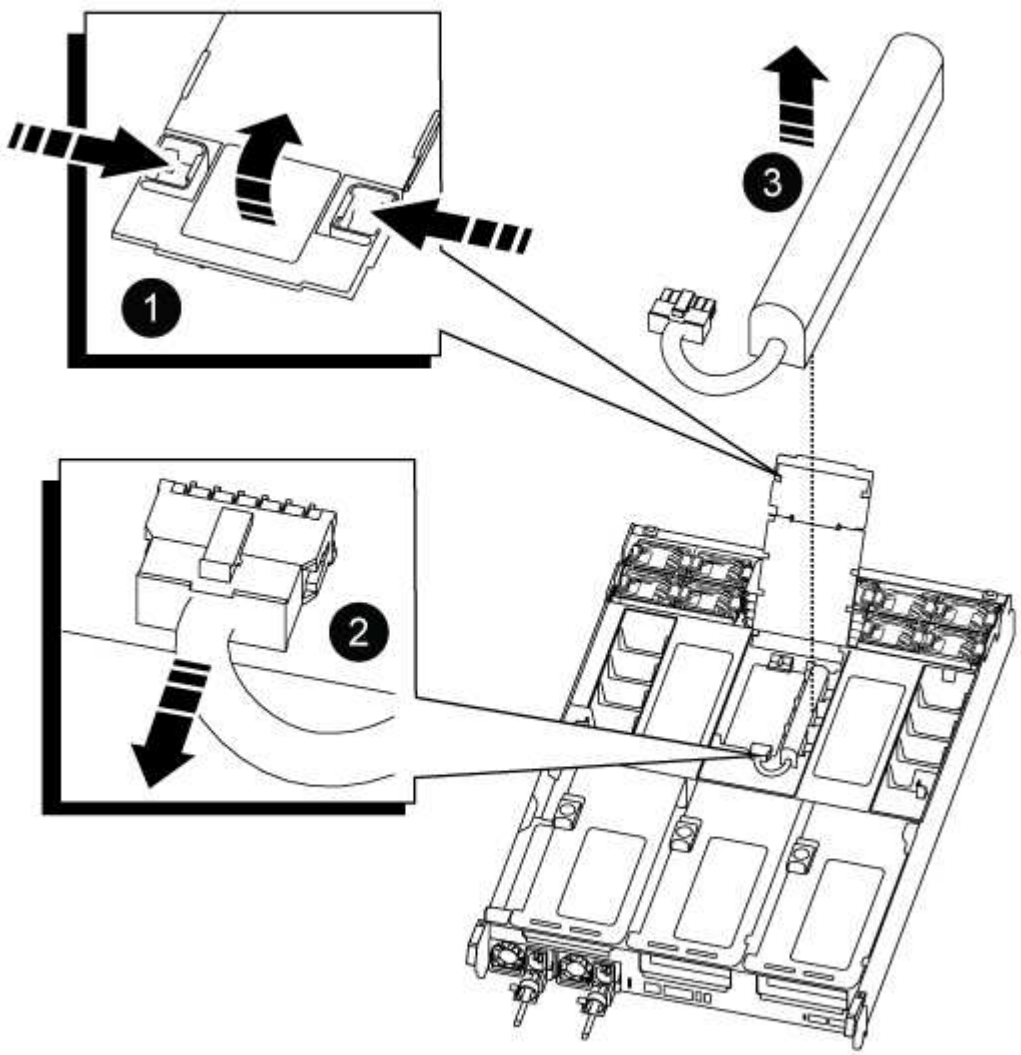

1	Air duct riser

2	NVDIMM battery plug
3	NVDIMM battery pack

Attention: The NVDIMM battery control board LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

1. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
2. Grasp the battery and lift the battery out of the air duct and controller module, and then set it aside.
3. Remove the replacement battery from its package.
4. Install the replacement battery pack in the NVDIMM air duct:
   a. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   b. Plug the battery plug into the riser socket and make sure that the plug locks into place.
5. Close the NVDIMM air duct.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

3. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe card - ASA C800

To replace a PCIe card, you must disconnect the cables from the cards, remove the SFP and QSFP modules from the cards before removing the riser, reinstall the riser, and then reinstall the SFP and QSFP modules before cabling the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

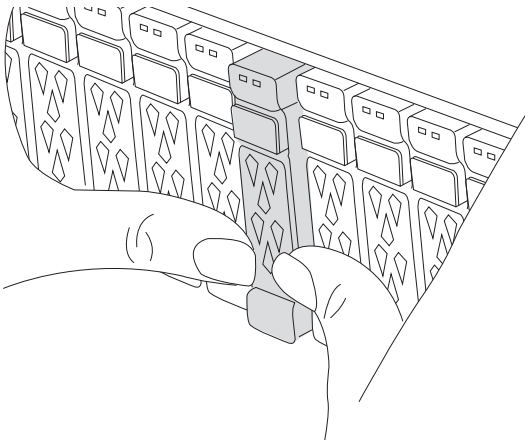

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

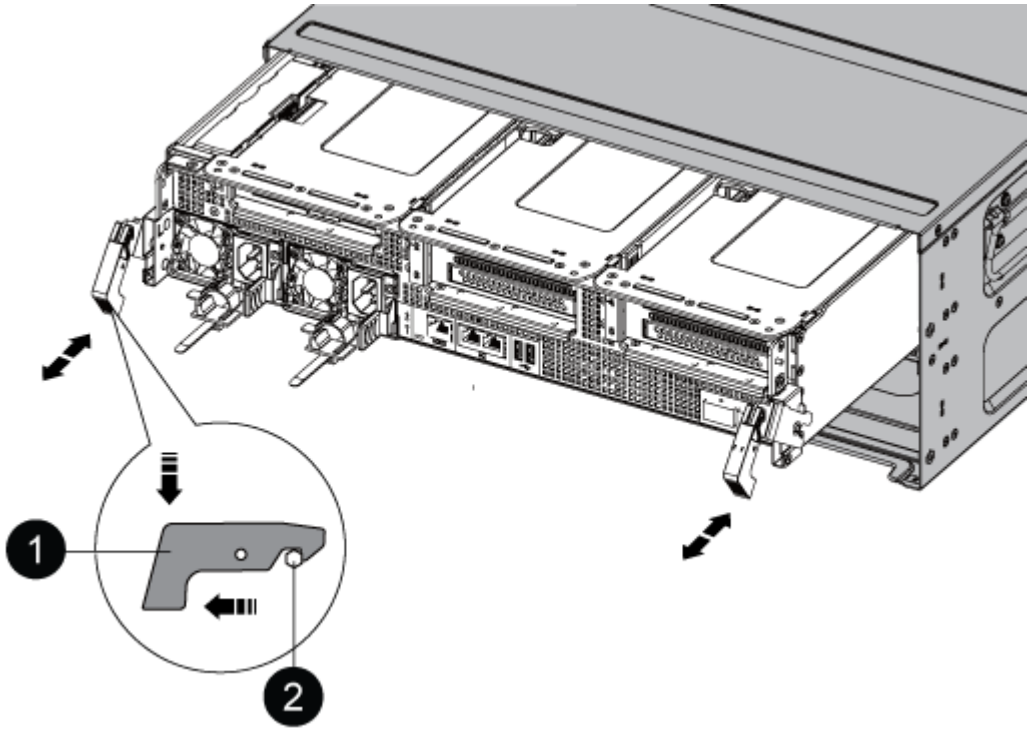

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

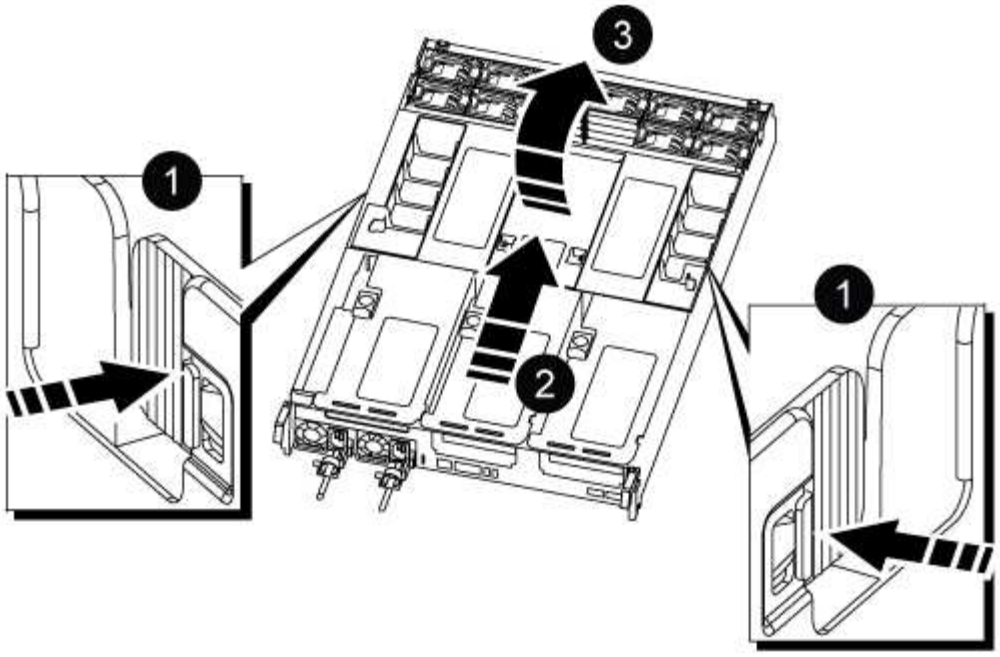

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace a PCIe card

To replace a PCIe card, you must remove the cabling and any QSFPs and SFPs from the ports on the PCIe cards in the target riser, remove the riser from the controller module, remove and replace the PCIe card, reinstall the riser and any QSFPs and SFPs onto the ports, and cable the ports.

1. Determine if the card you are replacing is from Riser 1 or if it is from Riser 2 or 3.
   - If you are replacing the 100GbE PCIe card in Riser 1, use Steps 2 - 3 and Steps 6 - 7.
   - If you are replacing a PCIe card from Riser 2 or 3, use Steps 4 through 7.
2. Remove Riser 1 from the controller module:
   a. Remove the QSFP modules that might be in the PCIe card.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

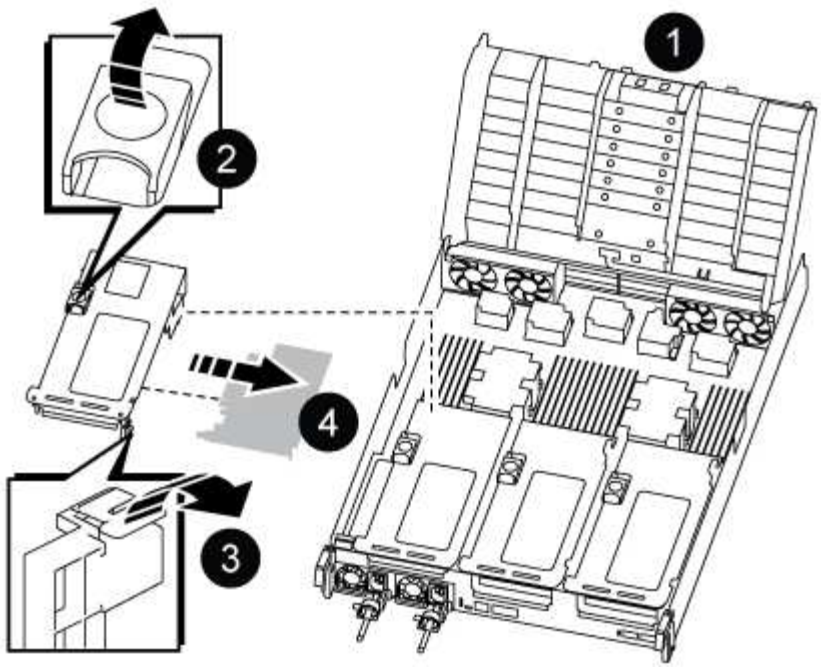

1	Air duct
2	Riser locking latch
3	Card locking bracket
4	Riser 1 (left riser) with 100GbE PCIe card in slot 1.

3. Remove the PCIe card from Riser 1:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. Remove the PCIe card from the riser.
4. Remove the PCIe riser from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

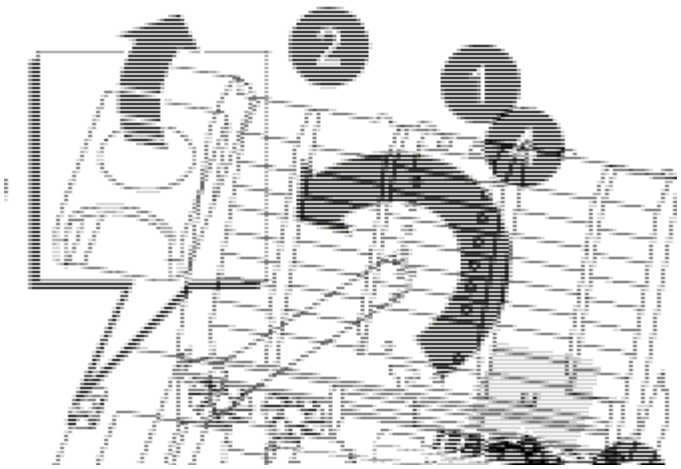

1	Air duct
2	Riser 2 (middle riser) or 3 (right riser) locking latch
3	Card locking bracket
4	Side panel on riser 2 or 3
5	PCIe cards in riser 2 or 3

5. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe cards.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. Swing the side panel off the riser.
   d. Remove the PCIe card from the riser.
6. Install the PCIe card into the same slot in the riser:
   a. Align the card with the card socket in the riser, and then slide it squarely into the socket in the riser.

      Make sure that the card is completely and squarely seated into the riser socket.

   b. For Riser 2 or 3, close the side panel.
   c. Swing the locking latch into place until it clicks into the locked position.
7. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the

controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

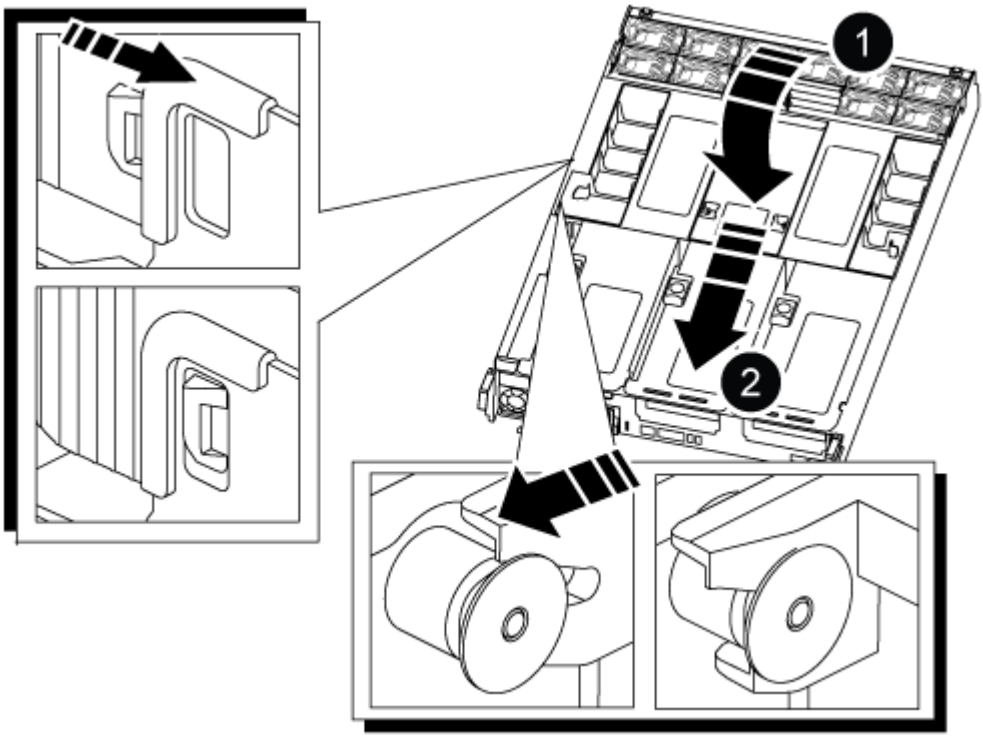

1	Locking tabs
2	Slide plunger

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
5. Plug the system cables and transceiver modules into the controller module and reinstall the cable management device.
6. Plug the power cables into the power supplies and reinstall the power cable retainers.

   The controller module begins to boot as soon as it is connected to power. Be prepared to interrupt the boot process.

   > If your system has DC power supplies, make sure the thumbscrews on the power supply cable are tight.

7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - ASA C800

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

About this task

This procedure is written for replacing one PSU at a time.

> It is a best practice to replace the PSU within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded PSU until the PSU is replaced.

> Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
   b. Unplug the power cable from the power source.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

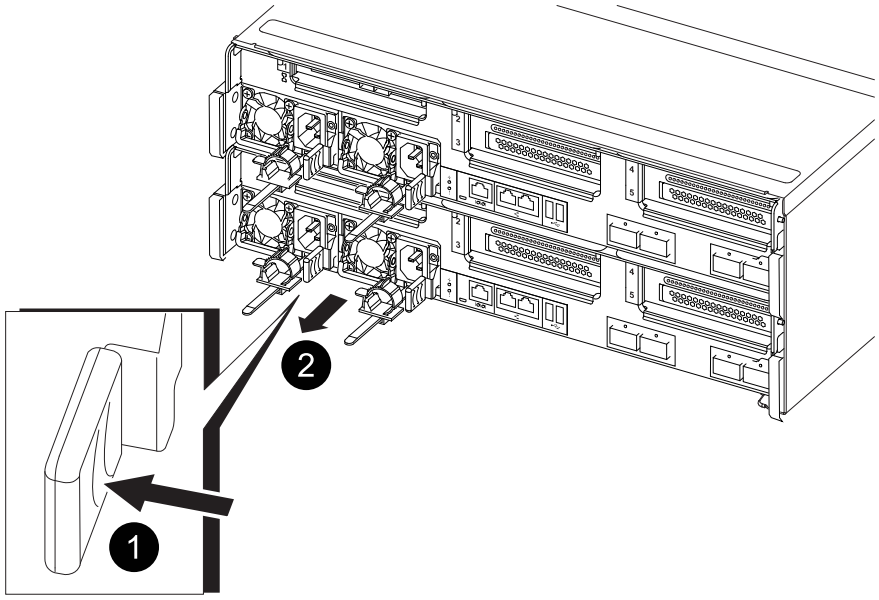

1	Blue PSU locking tab
2	Power supply

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU and the power source.
   b. Secure the power cable to the PSU using the power cable retainer.

   Once power is restored to the PSU, the status LED should be green.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

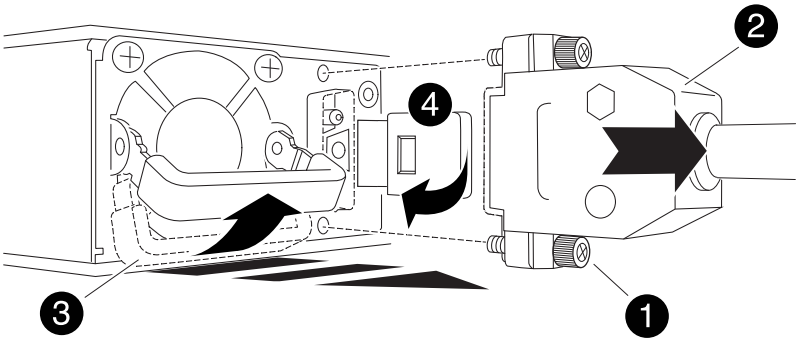

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle

4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA C800

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

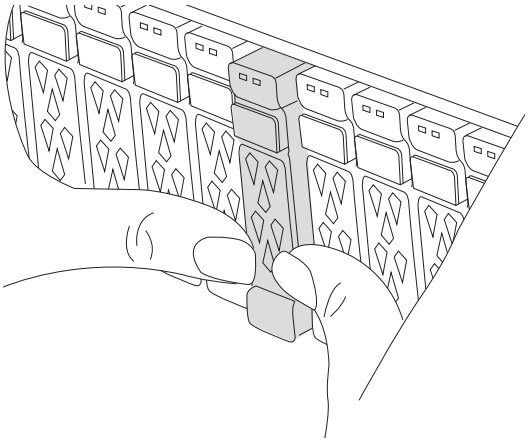

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
6. Remove the cable management device from the controller module and set it aside.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

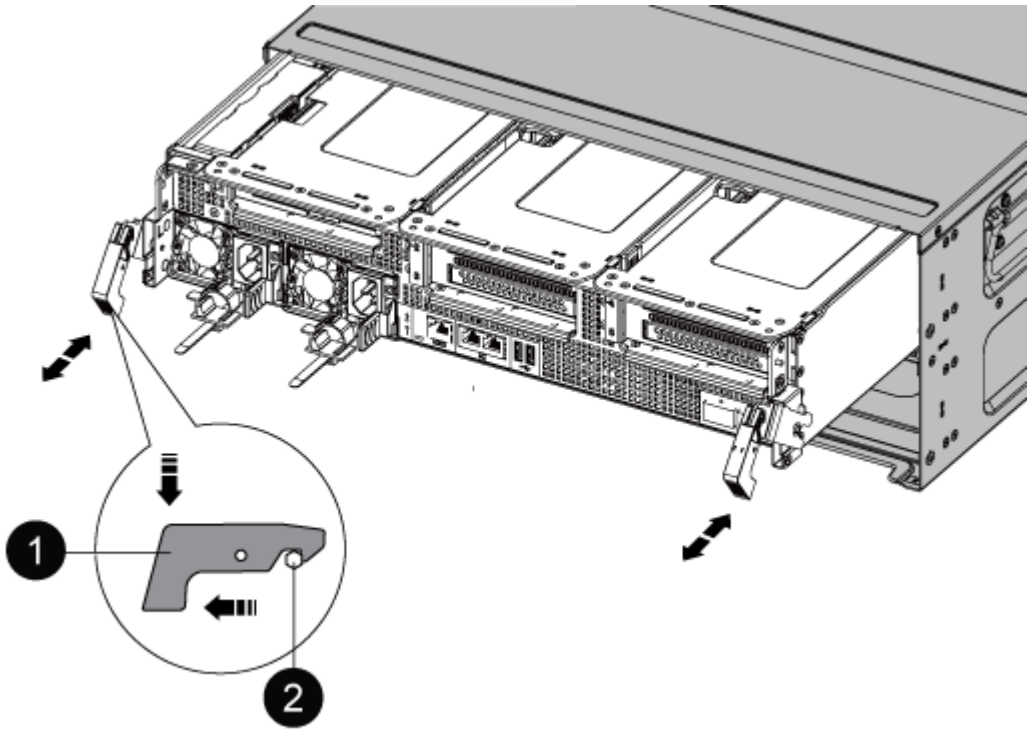

1	Locking latch
2	Locking pin

1. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

2. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

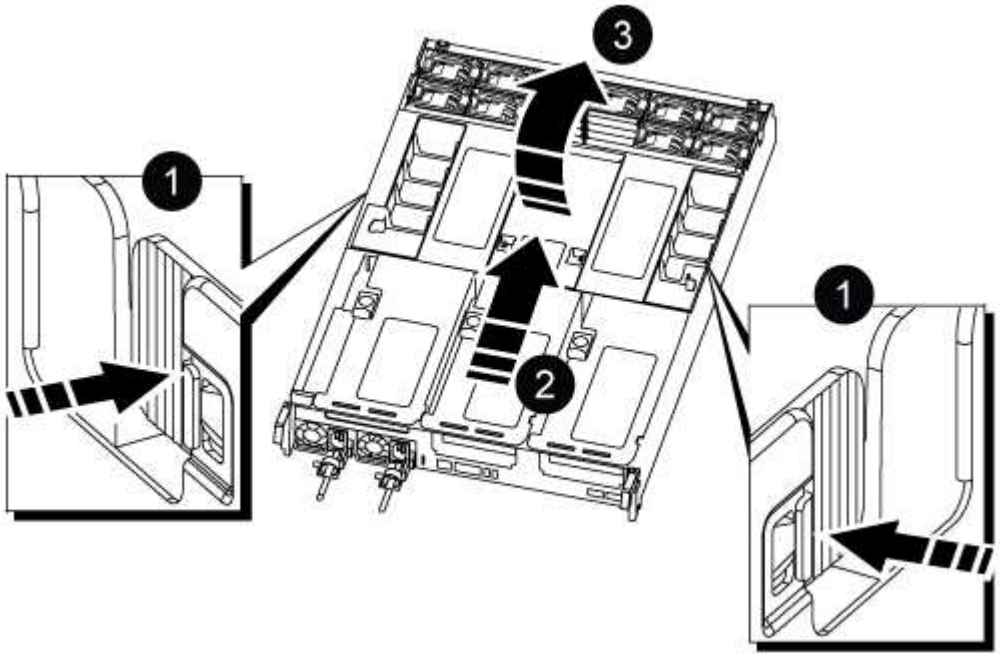

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 3: Replace the RTC battery

Original controller

1. Remove PCIe riser 2 (middle riser) from the controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

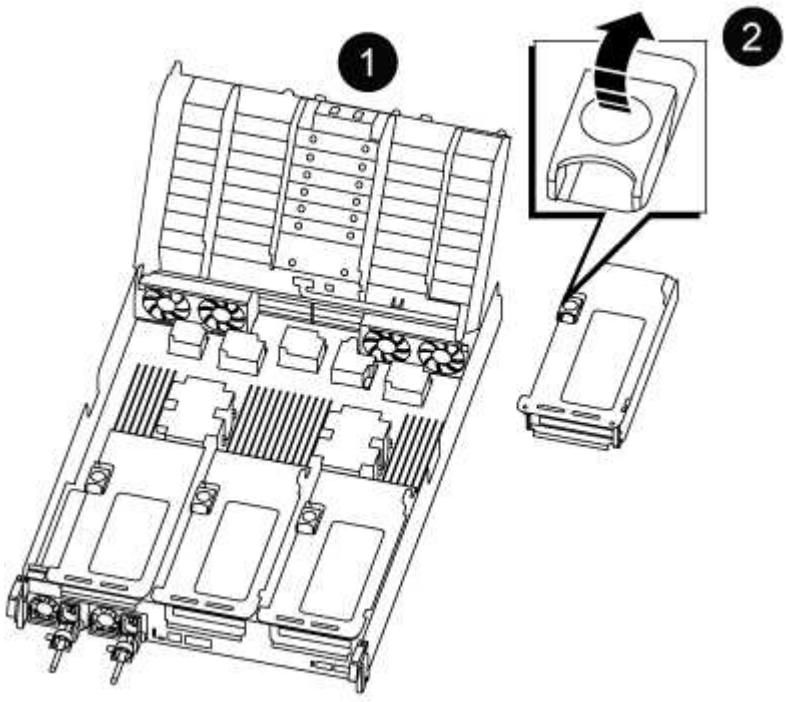

1	Air duct
2	Riser 2 (middle riser) locking latch

2. Locate the RTC battery under Riser 2.

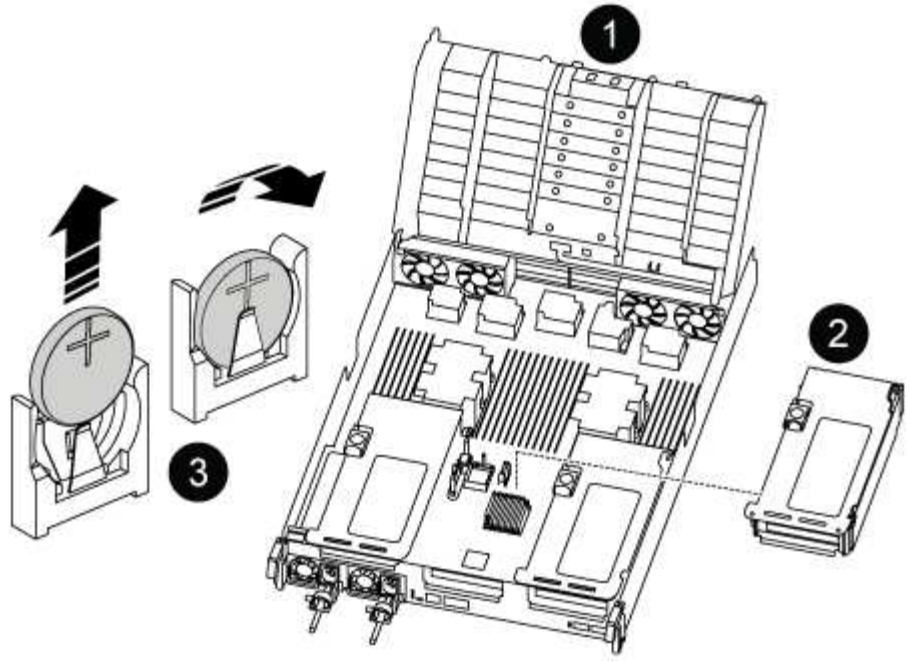

1	Air duct
2	Riser 2
3	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
7. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

VER2 controller

1. Locate the RTC battery near the DIMMs.

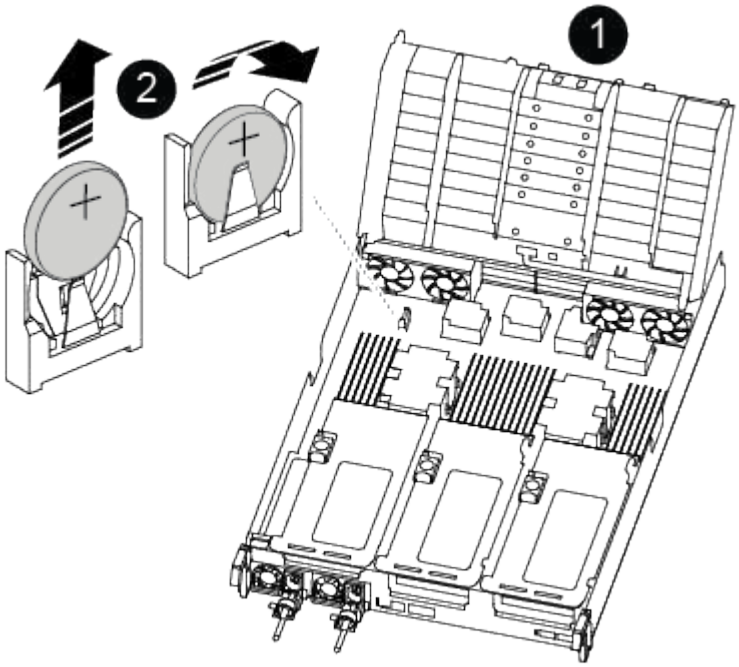

1	Air duct
2	RTC battery and housing

2. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

3. Remove the replacement battery from the antistatic shipping bag.
4. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
5. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Halt the controller at the LOADER prompt.
5. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA r2 systems

Install and setup your ASA r2 systems

Go to the ASA r2 install and setup instructions to learn how to install and setup your system.

The ASA r2 systems documentation includes information such as:

- Installation and setup instructions
- Administration instructions for configuring your system, such as provisioning SAN storage, cloning data, and re-sizing local storage.
- Instructions for managing your system, including managing client access, securing your data, and protecting your data.
- Monitoring and troubleshooting instructions, including information about alerts, cluster events, and system logs.

Specific maintenance procedures for each type of ASA r2 system are in the ASA r2 systems maintenance section.

Maintain ASA r2 systems

ASA A1K systems

Overview of the maintenance procedures - ASA A1K

Maintain the hardware of your ASA A1K storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the ASA A1K system has already been deployed as a storage node in the ONTAP environment.

System components

For the ASA A1K storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media.
Controller	A controller consists of a board, firmware, and software. It controls the storage, I/O cards, and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.

Fan	A fan cools the controller.
NVRAM	The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module.
NV battery	The NV battery is responsible for providing power to the NVRAM module while data in-flight is being destaged to flash memory after a power loss.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.
System management module	The System management module provides the interface between the controller and a console or laptop for controller or system maintenance purposes. The System management module contains the boot media and stores the system serial number (SSN).

Boot media

Boot media replacement workflow - ASA A1K

Get started with replacing the boot media in your ASA A1K storage system by reviewing the replacement requirements, shutting down the controller, replacing the boot media, restoring the image on the boot media, and verifying the system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the System Management module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements to replace the boot media - ASA A1K

Before replacing the boot media in your ASA A1K system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming the cluster ports on the impaired controller are working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

Review the following requirements.

- You must replace the failed boot media with a replacement boot media you received from NetApp.
- The cluster ports are used to communicate between the two controllers during the automated boot recovery process. Make sure that the cluster ports on the impaired controller are working properly.
- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg
  - /cfcard/kmip/certs/client.crt
  - /cfcard/kmip/certs/client.key
  - /cfcard/kmip/certs/CA.pem
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller to replace the boot media - ASA A1K

Shut down the impaired controller in your ASA A1K storage system to prevent data loss and ensure system stability when replacing the boot media.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media - ASA A1K

The boot media in your ASA A1K system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media in the System Management module, and then reinstalling the System Management module.

The boot media is located inside the System Management module and is accessed by removing the module from the system.

Replace the boot media.

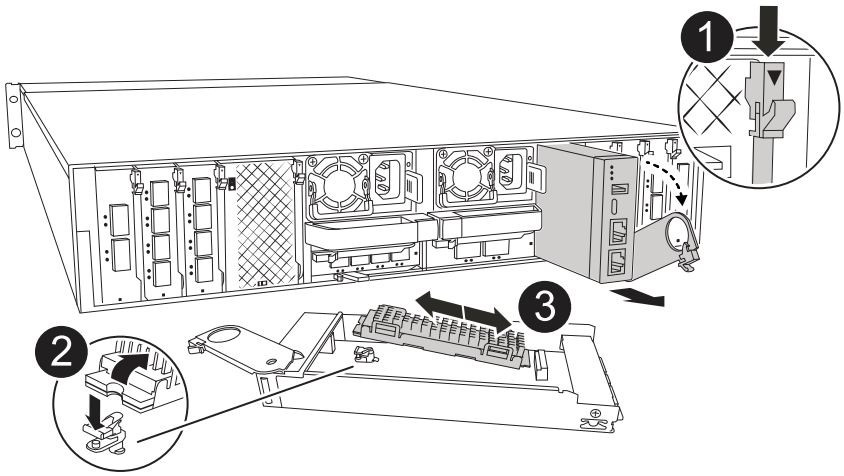

1	System Management module cam latch
2	Boot media locking button
3	Boot media

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.

   If your storage system has DC power supplies, disconnect the power cable block from the power supply units (PSUs).

3. Remove the System Management module:
   a. Remove any cables connected to the System Management module. Make sure that you label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   c. Depress the System Management cam button.
   d. Rotate the cam latch down as far as it will go.
   e. Remove the System Management module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
   f. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:
   a. Press the blue locking button.
   b. Rotate the boot media up, slide it out of the socket, and set it aside.
5. Install the replacement boot media into the System Management module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down toward the locking button.

c. Push the locking button, rotate the boot media all the way down and then release the locking button.

6. Reinstall the System Management module:

a. Align the module with the edges of the enclosure slot opening.

b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.

7. Rotate the cable management tray up to the closed position.

a. Recable the System Management module.

8. Plug the power cables into the power supplies and reinstall the power cable retainer.

The controller begins to boot as soon as power is reconnected to the system.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Restore the ONTAP image on the boot media - ASA A1K

After installing the new boot media device in your ASA A1K system, you can start the automated boot media recovery process to restore the configuration from the partner node.

During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

```
boot_recovery -partner
```

The screen displays the following message:

```
Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
```

2. Monitor the boot media install recovery process.

The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed part to NetApp - ASA A1K

If a component in your ASA A1K system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - ASA A1K

Get started with replacing the controller in your ASA A1K storage system by shutting down the impaired controller, removing and replacing the controller, reconfiguring the system settings, and verifying system operations.

Review controller replacement requirements

To replace the controller module, you must meet certain requirements.

Shut down the impaired controller

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

Replace the controller

Replacing the controller includes removing the impaired controller, moving the FRU components to the replacement controller module, and then installing the replacement controller module in the enclosure.

Restore and verify the system configuration

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Give back the controller

Transfer the ownership of storage resources back to the replacement controller.

Complete controller replacement

Verify the LIfs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - ASA A1K

Before replacing the controller in your ASA A1K system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- You must replace the failed component with a replacement component you received from NetApp.
- Do not use this procedure for controller upgrades; instead, refer to the Choose your controller hardware upgrade procedure for guidance.
- You must replace a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text log file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

What's next?

After you've reviewed the requirements to replace your ASA A1K controller, you need to shut down the controllers.

Shut down the impaired controller - ASA A1K

Shut down the controller in your ASA A1K storage system to prevent data loss and ensure system stability when replacing the controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be

resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After you've shut down the controller, you need to replace the controller.

Replace the controller - ASA A1K

Replace the controller in your ASA A1K system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

Step 1: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

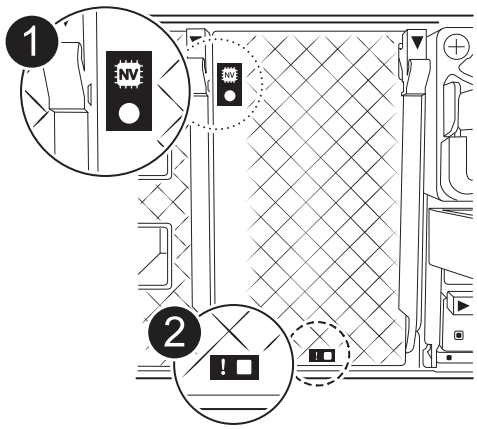

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

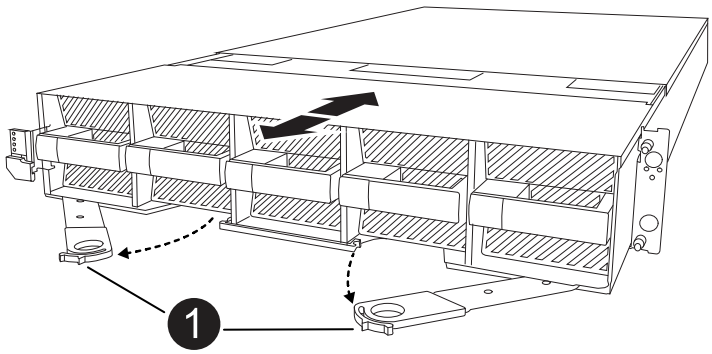

1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 2: Move the fans

You must remove the five fan modules from the impaired controller module to the replacement controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press the gray locking button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

> The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

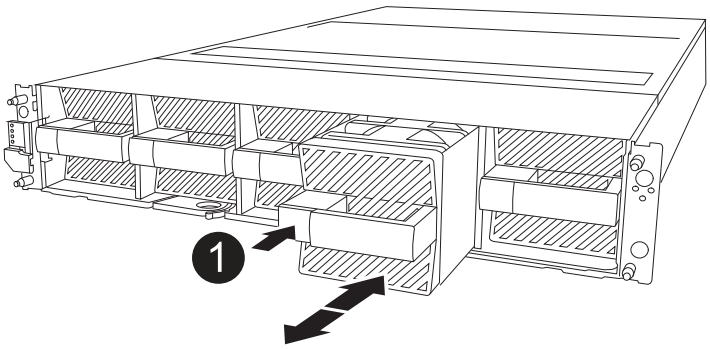

1	Black locking button

4. Install the fan in the replacement controller module:
   a. Align the edges of the fan housing with the opening in the front of the replacement controller module.
   b. Gently slide the fan module all the way into the replacement controller module until it locks in place.
5. Repeat the preceding steps for the remaining fan modules.

Step 3: Move the NV battery

Move the NV battery to the replacement controller.

Steps

1. Open the NV battery air duct cover and locate the NV battery.

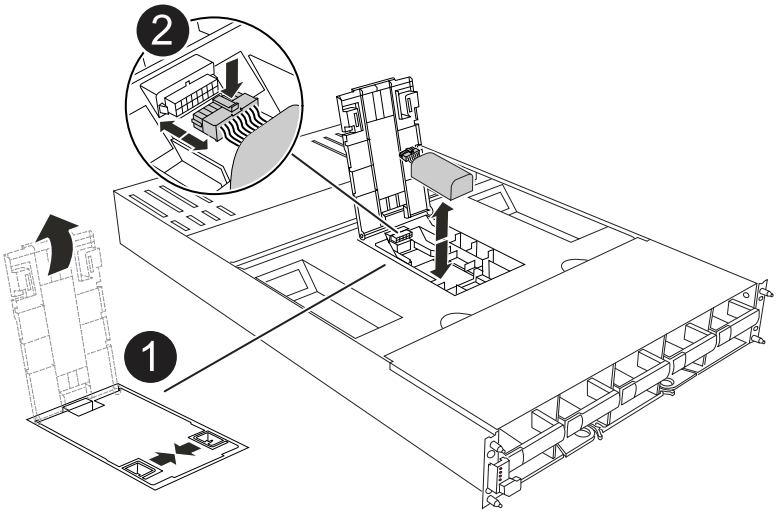

1	NV battery air duct cover
2	NV battery plug
3	NV battery pack

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module.
5. Move the battery pack to the replacement controller module and then install it in the NV battery air duct:
   a. Open the NV battery air duct in the replacement controller module.
   b. Plug the battery plug into the socket and make sure that the plug locks into place.
   c. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   d. Close the air duct cover.

Step 4: Move system DIMMs

Move the DIMMs to the replacement controller module.

Steps

1. Open the motherboard air duct and locate the DIMMs.

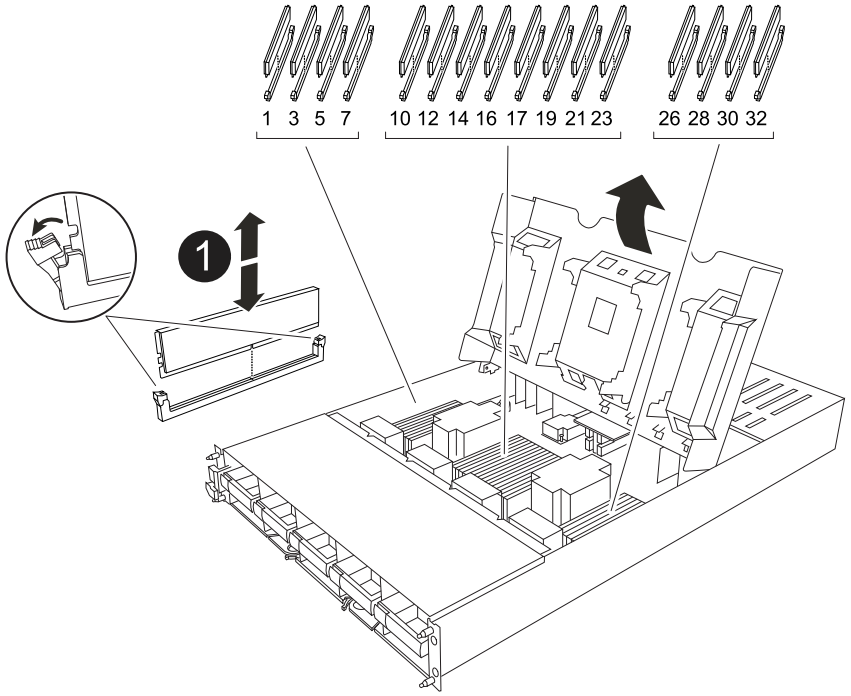

1	System DIMM

2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

4. Locate the slot where you are installing the DIMM in the replacement controller module.
5. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Repeat these steps for the remaining DIMMs.
   Close the motherboard air duct.

Step 5: Install the controller module

Reinstall the controller module and boot it.

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

   Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

   The controller boots to the LOADER prompt as soon as it is fully seated.

4. From the LOADER prompt, enter `show date` to display the date and time on the replacement controller. Date and time are in GMT.

   Time displayed is local time not always GMT and is displayed in 24hr mode.

5. Set the current time in GMT with the `set time hh:mm:ss` command. You can get the current GMT from the partner node the `date -u`command.
6. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

What's next?

After you've replaced the impaired ASA A1K controller, you need to restore the system configuration.

Restore and verify the system configuration - ASA A1K

Verify that the controller's HA configuration is active and functioning correctly in your ASA A1K storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

Steps

1. Boot to maintenance mode: `boot_ontap maint`
   a. Enter `y` when you see *Continue with boot?*.

      If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state: `ha-config show`

The HA state should be the same for all components.

5. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha`

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed: `ha-config show`

Step 2: Verify disk list

Steps

1. Verify that the adapter lists the paths to all disks with the `storage show disk -p`.

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode: `halt`.

What's next?

After you've restored and verified the system configuration for your ASA A1K system, you need to give back the controller.

Give back the controller - ASA A1K

Return control of storage resources to the replacement controller so your ASA A1K system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`
   - If you encounter errors, contact NetApp Support.
9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

    c. Verify status of the keys: `security key-manager key query -restored false`

The output should show no results when when properly synchronized.

If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

After you've transferred the ownership of storage resources back to the replacement controller, you need to complete the controller replacement procedure.

Complete controller replacement - ASA A1K

To complete the controller replacement for your ASA A1K system, first restore the NetApp Storage Encryption configuration (if necessary). Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, return the failed part to NetApp.

Step 1: Verify LIFs and check cluster health

Before returning the replacement node to service, verify that the logical interfaces are on their home ports, check the cluster health, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any logical interfaces are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - ASA A1K

Replace a DIMM in your ASA A1K system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting

ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- Make sure all other components in the system are functioning properly; if not, you must contact technical support.
- Make sure you replace the failed component with a replacement component you received from NetApp.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

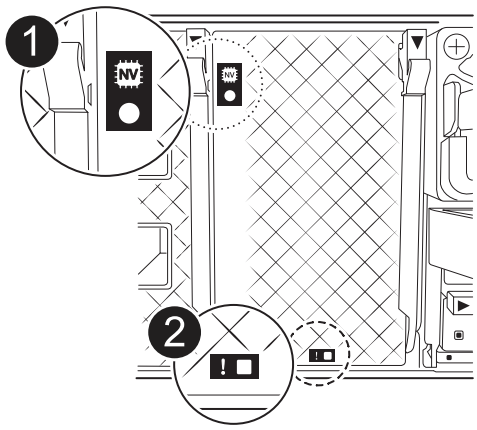

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

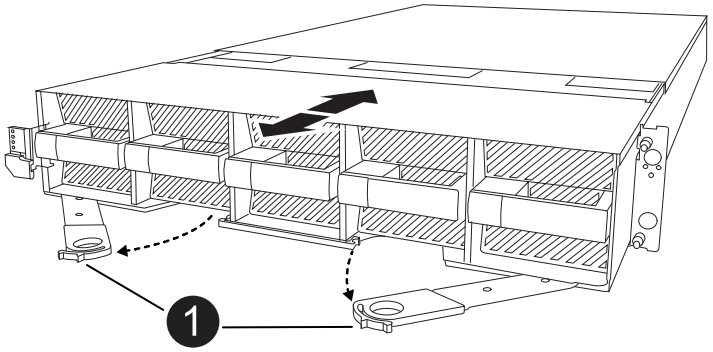

1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace a DIMM

You must replace a DIMM when the system reports a permanent failure condition for that DIMM.

Steps

1. If you are not already grounded, properly ground yourself.
2. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
3. Locate the DIMMs on your controller module and identify the DIMM for replacement.

   Use the FRU map on the controller airduct to locate the DIMM slot.

4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

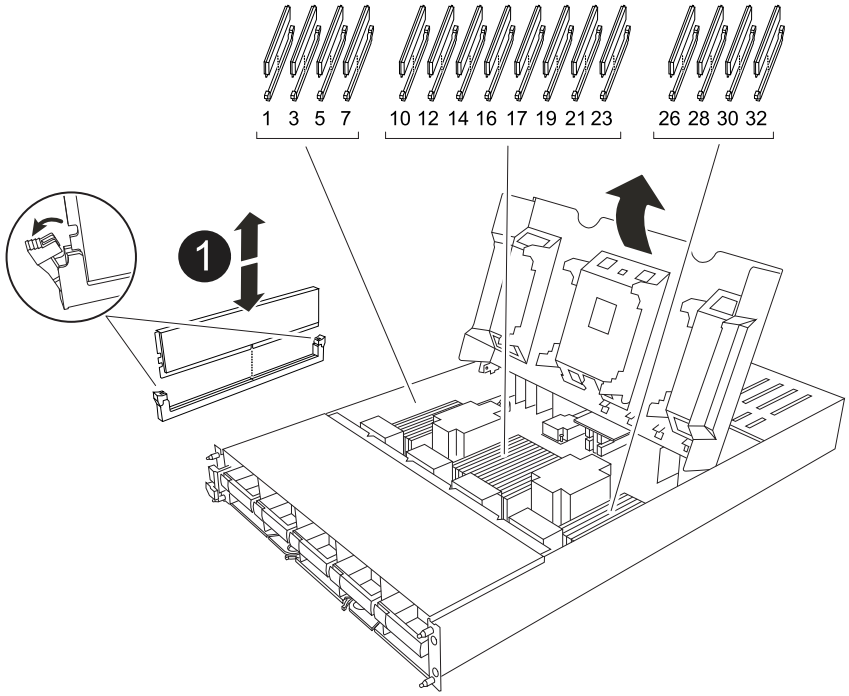

1	DIMM and DIMM ejector tabs

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the controller air duct.

Step 4: Install the controller

Reinstall the controller module and boot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch

back under the fans

> Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

The controller module begins to boot as soon as it is fully seated in the enclosure.

4. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
6. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a fan - ASA A1K

Replace a fan module in your ASA A1K system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.

   Facing the controller module, fan modules are numbered 1 through 5, from left to right.

   > There is a single LED for each fan. It is green when the fan is functioning correctly and amber when not.

4. Press the black button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

   > The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

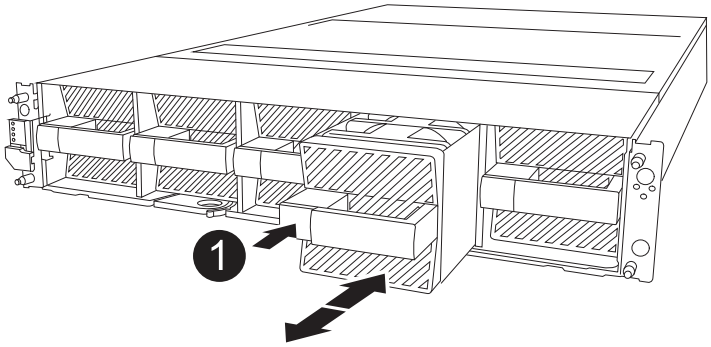

1	Black release button

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED turns off once the fan is recognized by that system.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace NVRAM - ASA A1K

Replace the NVRAM in your ASA A1K system when the non-volatile memory becomes faulty or requires an upgrade. The replacement process involves shutting down the impaired controller, replacing the NVRAM module or the NVRAM DIMM, reassigning the disks, and returning the failed part to NetApp.

The NVRAM module consists of the NVRAM12 hardware and field-replaceable DIMMs. You can replace a failed NVRAM module or the DIMMs inside the NVRAM module.

Before you begin

- Make sure you have the replacement part available. You must replace the failed component with a replacement component you received from NetApp.
- Make sure all other components in the storage system are functioning properly; if not, contact NetApp support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv

advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the NVRAM module or NVRAM DIMM

Replace the NVRAM module or NVRAM DIMMs using the appropriate following option.

Option 1: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 4/5 in the enclosure and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
4. Remove the impaired NVRAM module from the enclosure:
   a. Depress the locking cam button.
   b. Rotate the cam latch down as far as it will go.
   c. Remove the impaired NVRAM module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

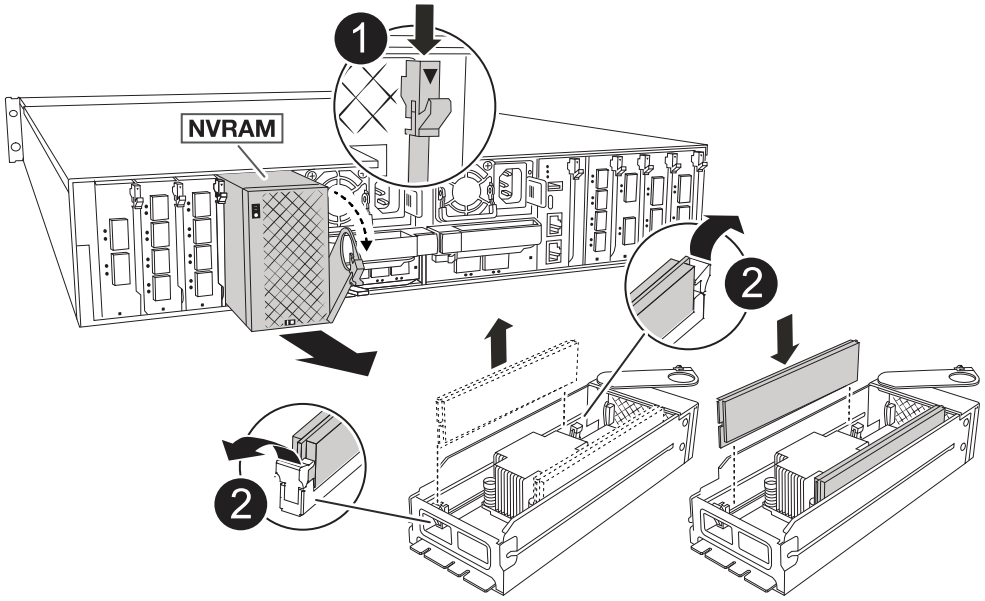

1	Cam locking button
2	DIMM locking tabs

5. Set the NVRAM module on a stable surface.
6. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
7. Install the replacement NVRAM module into the enclosure:
   a. Align the module with the edges of the enclosure opening in slot 4/5.
   b. Gently slide the module into the slot all the way, and then rotate the cam latch all the way up to lock the module in place.

8. Recable the PSUs.
9. Rotate the cable management tray up to the closed position.

Option 2: Replace the NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, you must remove the NVRAM module, and then replace the target DIMM.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the controller PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
4. Remove the target NVRAM module from the enclosure.

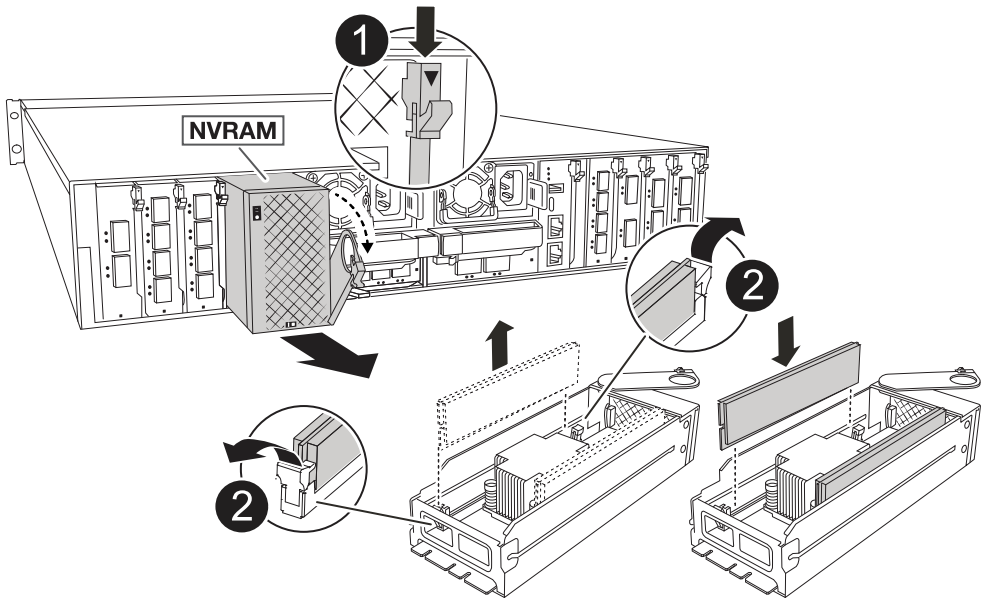

1	Cam locking button
2	DIMM locking tabs

5. Set the NVRAM module on a stable surface.
6. Locate the DIMM to be replaced inside the NVRAM module.

> Consult the FRU map label on the side of the NVRAM module to determine the locations of DIMM slots 1 and 2.

7. Remove the DIMM by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
8. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
9. Install the NVRAM module into the enclosure:

a. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

10. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace the component, you must reboot the controller module by plugging the power cables back into the PSU.

The system will begin to reboot, typically to the LOADER prompt.

Step 4: Verify controller state

You must confirm the controller state of the controllers connected to the disk pool when you boot the controller.

Steps

1. If the controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: *halt*
2. From the LOADER prompt on the controller, boot the controller and enter *y* when prompted to override the system ID due to a system ID mismatch.
3. Wait until the Waiting for giveback… message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify the system state: *storage failover show*

   In the command output, you should see a message indicates the state of the controllers.

```
                              Takeover
Node           Partner        Possible State Description
-------------- -------------- --------
-------------------------------------
<nodename>
               <nodename>-   true     Connected to <nodename>-P2-3-178.
               P2-3-178                Waiting for cluster applications
to
                                       come online on the local node.
AFF-A90-NBC-P2-3-178
               <nodename>-   true     Connected to <nodename>-P2-3-177,
               P2-3-177                Partial giveback
2 entries were displayed.
```

4. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: *storage failover giveback -ofnode replacement_node_name*

      The controller connects back its storage pool and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter *y*.

If the giveback is vetoed, you can consider overriding the vetoes.

For more information, see the Manual giveback commands topic to override the veto.

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: *storage failover show*

5. Verify all disks are displayed: `storage disk show`

```
::> storage disk show
                     Usable           Disk    Container   Container
Disk                   Size Shelf Bay Type    Type        Name
---------------- ---------- ----- --- ------- ----------- ---------
1.0.0                3.49TB     0   0 SSD-NVM aggregate   pod_NVME_SSD_1
1.0.1                3.49TB     0   1 SSD-NVM aggregate   pod_NVME_SSD_1
1.0.2                3.49TB     0   2 SSD-NVM aggregate   pod_NVME_SSD_1
1.0.3                3.49TB     0   3 SSD-NVM aggregate   pod_NVME_SSD_1
1.0.4                3.49TB     0   4 SSD-NVM aggregate   pod_NVME_SSD_1

[...]
48 entries were displayed.
```

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NV battery - ASA A1K

Replace the NV battery in your ASA A1K system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

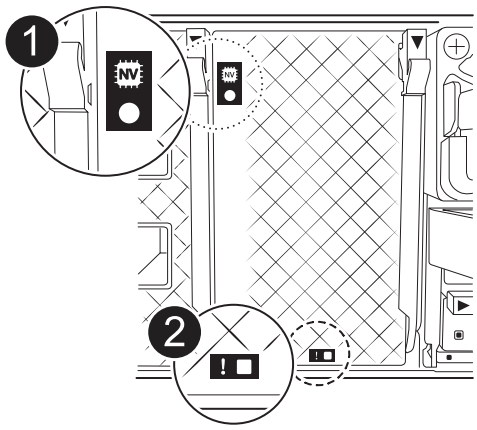

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

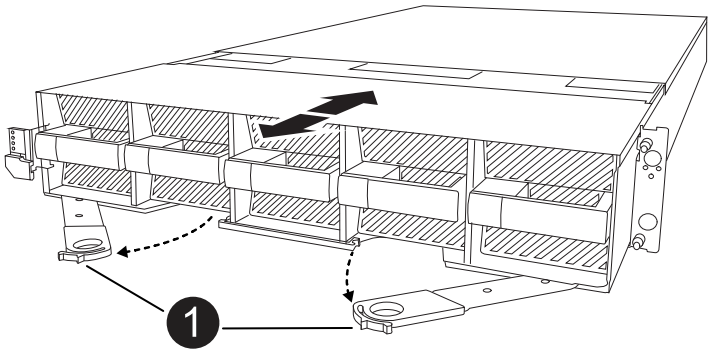

1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace the NV battery

Remove the failed NV battery from the controller module and install the replacement NV battery.

Steps

1. Open the air duct cover and locate the NV battery.

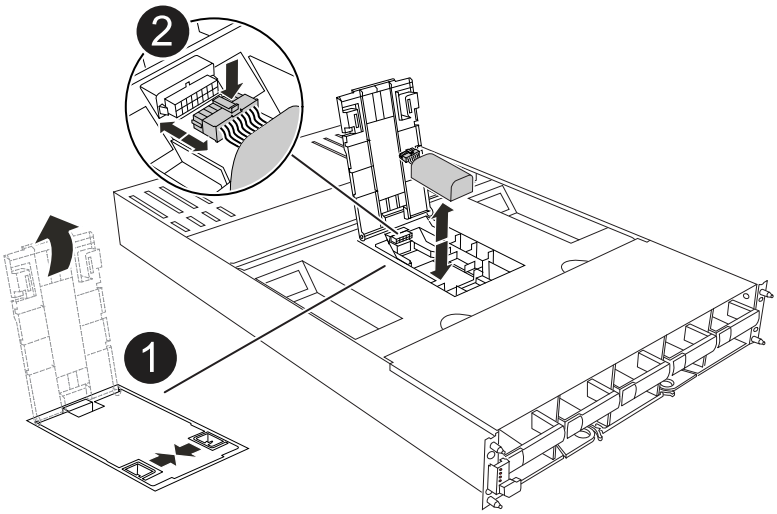

1	NV battery air duct cover
2	NV battery plug

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module, and then set it aside.
5. Remove the replacement battery from its package.
6. Install the replacement battery pack into the controller:
   a. Plug the battery plug into the riser socket and make sure that the plug locks into place.
   b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
7. Close the NV air duct cover.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.

3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

   > Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the enclosure.

4. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
6. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Overview of add and replace an I/O module - ASA A1K

The ASA A1K system offers flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your ASA A1K storage system with the same type of I/O module, or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

- Add an I/O module

  Adding additional modules can improve redundancy, helping to ensure that the system remains operational even if one module fails.

- Replace an I/O module

  Replacing a failing I/O module can restore the system to its optimal operating state.

I/O slot numbering

The I/O slots on ASA A1K controllers are numbered 1 through 11, as shown in the following illustration.

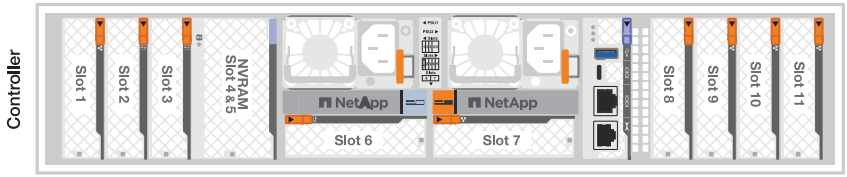

Add an I/O module to your ASA A1K system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your ASA A1K storage system when there are empty slots available or when all slots are fully populated.
.About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

A storage system has two location LEDs, one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- Make sure that all other components are functioning properly.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the target slot blanking module from the carrier:
   a. Depress the cam latch on the blanking module in the target slot.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
4. Install the I/O module:
   a. Align the I/O module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
5. Cable the I/O module to the designated device.

   Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

6. Rotate the cable management tray up to the closed position.
7. From the LOADER prompt, reboot the node:

   ```
   bye
   ```

   This reinitializes the I/O module and other components and reboots the node.

8. Give back the controller from the partner controller:

   ```
   storage failover giveback -ofnode target_node_name
   ```

9. Repeat these steps for controller B.
10. From the healthy node, restore automatic giveback if you disabled it:

    ```
    storage failover modify -node local -auto-giveback true
    ```

11. If AutoSupport is enabled, restore automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the target I/O module from the chassis:
   a. Depress the cam latch button.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

      Make sure that you keep track of which slot the I/O module was in.
5. Install the I/O module into the target slot in the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
6. Cable the I/O module to the designated device.
7. Repeat the remove and install steps to replace additional modules for the controller.
8. Rotate the cable management tray up to the closed position.
9. Reboot the controller from the LOADER prompt:_bye_

   This reinitializes the PCIe cards and other components and reboots the node.

   If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement
10. Give back the controller from the partner controller:

    ```
    storage failover giveback -ofnode target_node_name
    ```
11. Enable automatic giveback if it was disabled:

```
storage failover modify -node local -auto-giveback true
```

12. Do one of the following:
    - If you removed a NIC I/O module and installed a new NIC I/O module, use the following network command for each port:

      ```
      storage port modify -node *<node name> -port *<port name> -mode network
      ```

    - If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
13. Repeat these steps for controller B.

Replace an I/O module - ASA A1K

Replace an I/O module in your ASA A1K system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

You can use this procedure with all versions of ONTAP supported by your storage system.

Before you begin

- You must have the replacement part available.
- Make sure all other components in the storage system are functioning properly; if not, contact technical support.

Step 1: Shut down the impaired node

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace a failed I/O module

To replace an I/O module, locate it within the enclosure and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

This following illustration shows removing a horizontal and vertical I/O module. Typically, you will only remove one I/O module.

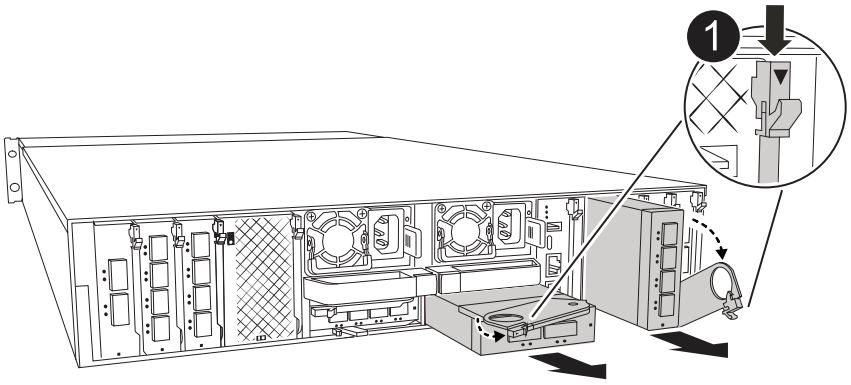

1	I/O cam latch

Make sure that you label the cables so that you know where they came from.

4. Remove the target I/O module from the enclosure:
   a. Depress the cam button on the target module.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

      Make sure that you keep track of which slot the I/O module was in.
5. Set the I/O module aside.
6. Install the replacement I/O module into the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
7. Cable the I/O module.
8. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```

3. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

4. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - ASA A1K

Replace an AC or DC power supply unit (PSU) in your ASA A1K system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cable, replacing the faulty PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

About this task

This procedure is written for replacing one PSU at a time.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

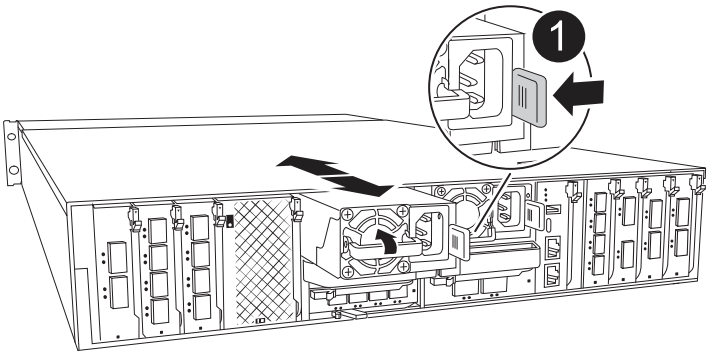

1	Terracotta PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU.
   b. Secure the power cable to the PSU using the power cable retainer.

   Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

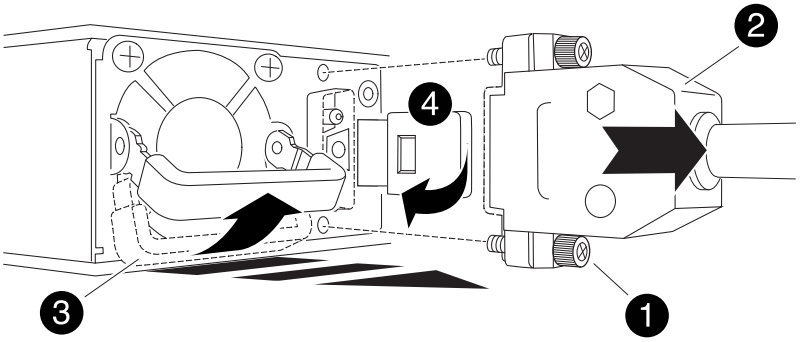

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle
4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one

way.

> To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA A1K

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your ASA A1K system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

- Understand that you can use this procedure with all versions of ONTAP supported by your system.
- Make sure all other components in the system are functioning properly; if not, you must contact contact NetApp support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

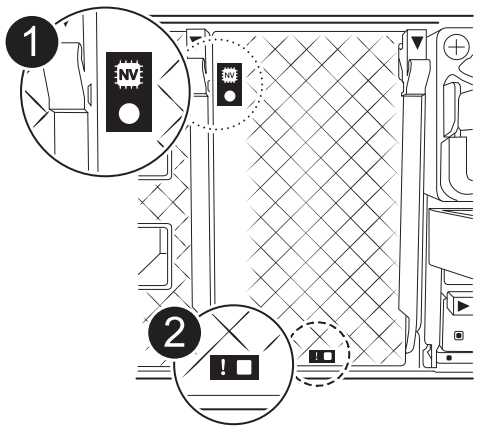

1	NVRAM status LED

2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

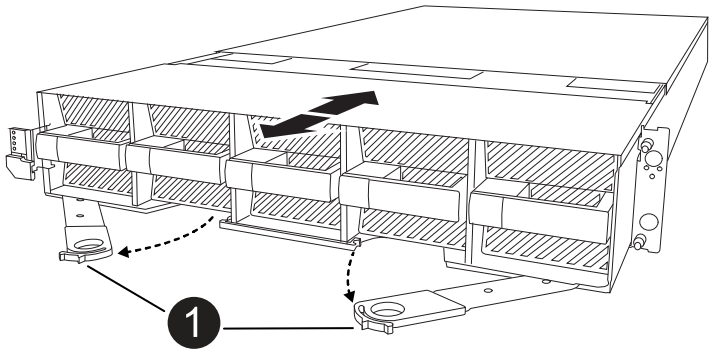

1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace the RTC battery

Remove failed RTC battery and install the replacement RTC battery.

Steps

1. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the RTC battery under the air duct.

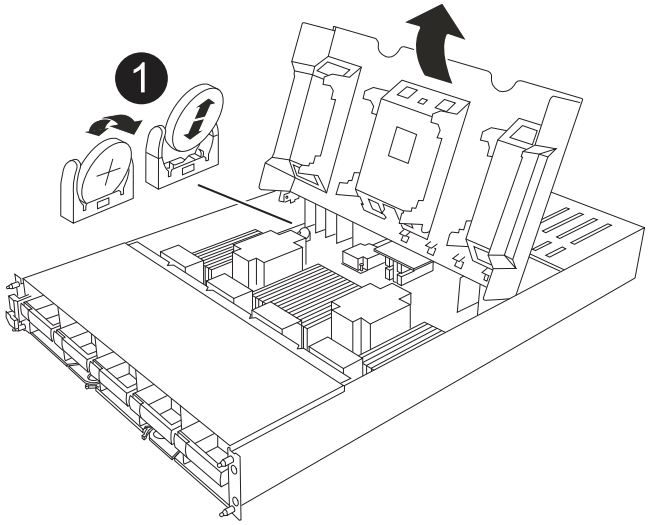

1	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

> Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the enclosure.

4. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
6. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Reset the time and date on the controller

> After replacing the RTC battery, inserting controller and powering on first BIOS reset, you will see the following error messages:
> `RTC date/time error. Reset date/time to default`
> `RTC power failure error`
> These messages are excpected and you can continue with this procedure.

Steps

1. Check the date and time on the healthy controller with the *cluster date show* command.

> If your system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*

a. At the LOADER prompt on the target controller, check the time and date with the `cluster date show` command.
b. If necessary, modify the date with the `set date mm/dd/yyyy` command.
c. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
    1. Confirm the date and time on the target controller.
    2. At the LOADER prompt, enter *bye* to reinitialize the PCIe cards and other components and let the controller reboot.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace system management module - ASA A1K

Replace the System Management module in your ASA A1K system when it becomes defective or its firmware is corrupted. The replacement process involves shutting down the controller, replacing the failed System Management module, rebooting the controller, updating the license keys, and returning the failed part to NetApp.

The System Management module, located at the back of the controller in slot 8, contains onboard components for system management, as well as ports for external management. The target controller must be shut down to replace an impaired System Management module or replace the boot media.

The System Management module has the following onboard components:

- Boot media, allowing boot media replacement without removing the controller module.

- BMC
- Management switch

The System Management module also contains the following ports for external management:

- RJ45 Serial
- USB Serial (Type-C)
- USB Type-A (Boot recovery)
- e0M RJ45 Ethernet

Before you begin

- All other system components must be working properly.
- The partner controller must be able to take over the impaired controller.
- You must replace the failed component with a replacement FRU component you received from your provider.

About this task

This procedure uses the following terminology:

- The impaired controller is the controller on which you are performing maintenance.
- The healthy controller is the HA partner of the impaired controller.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the impaired System Management module

Replace the impaired system management module.

Steps

1. Remove the System Management module:

   Make sure NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged. If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

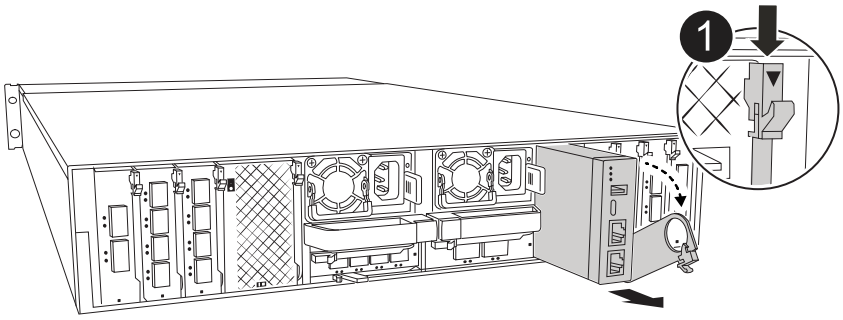

1	System Management module cam latch

   a. If you are not already grounded, properly ground yourself.
   b. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the

module.

c. Unplug the power supply cables from the PSUs.

d. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

e. Depress the cam button on the System Management module.

f. Rotate the cam lever down as far as it will go.

g. Loop your finger into the hole on the cam lever and pull the module straight out of the system.

h. Place the System Management module on an anti-static mat, so that the boot media is accessible.

2. Move the boot media to the replacement System Management module:

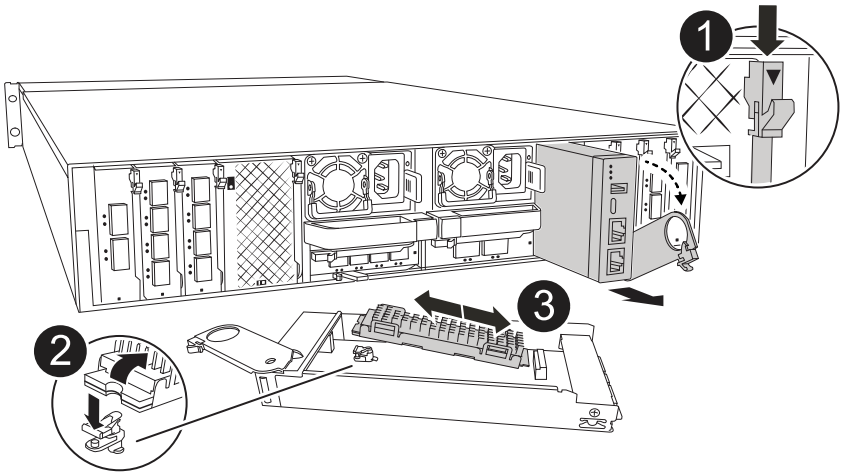

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue boot media locking button in the impaired System Management module.

b. Rotate the boot media up and slide it out of the socket.

3. Install the boot media in the replacement System Management module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down until it touches the locking button.

c. Depress the blue locking and rotate the boot media all the way down and release the blue locking button.

4. Install the replacement System Management module into the enclosure:

a. Align the edges of the replacement System Management module with the system opening and gently push it into the controller module.

b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

5. Rotate the cable management arm up to the closed position.
6. Recable the System Management module.

Step 3: Reboot the controller module

Reboot the controller module.

Steps

1. Plug the power cables back into the PSU.

   The system will begin to reboot, typically to the LOADER prompt.

2. Enter *bye* at the LOADER prompt.
3. Return the controller to normal operation by giving back its storage: *storage failover giveback -ofnode _impaired_node_name_*
4. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
5. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Step 4: Install licenses and register serial number

You must install new licenses for the node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the node. However, if the node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the for the node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

> The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`
4. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA A70 and ASA A90 systems

Overview of the maintenance procedures - ASA A70 and ASA A90

Maintain the hardware of your ASA A70 and ASA A90 storage systems to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the ASA A70 and ASA A90 systems have already been deployed as a storage node in the ONTAP environment.

System components

For the ASA A70 and ASA A90 storage systems, you can perform maintenance procedures on the following components.

Boot media	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the storage, I/O cards, and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.

Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller.
NVRAM	The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module.
NV battery	The NV battery is responsible for providing power to the NVRAM module while data in-flight is being destaged to flash memory after a power loss.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.
System Management module	The System Management module provides the interface between the controller and a console or laptop for controller or system maintenance purposes. The System management module contains the boot media and stores the system serial number (SSN).

Boot media

Boot media replacement workflow - ASA A70 and ASA A90

Get started with replacing the boot media in your ASA A70 and ASA A90 storage system by reviewing the replacement requirements, shutting down the controller, replacing the boot media, restoring the image on the boot media, and verifying the system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the System Management module and install the replacement boot media.

4 Restore the image on the boot media

Restore the ONTAP image from the partner controller.

5 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements to replace the boot media - ASA A70 and ASA A90

Before replacing the boot media in your ASA A70 or ASA A90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming the cluster ports on the impaired controller are working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

Review the following requirements.

- You must replace the failed boot media with a replacement boot media you received from NetApp.
- The cluster ports are used to communicate between the two controllers during the automated boot recovery process. Make sure that the cluster ports on the impaired controller are working properly.
- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg
  - /cfcard/kmip/certs/client.crt
  - /cfcard/kmip/certs/client.key
  - /cfcard/kmip/certs/CA.pem
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller to replace the boot media - ASA A70 or ASA A90

Shut down the impaired controller in your ASA A70 or ASA A90 storage system to prevent data loss and ensure system stability when replacing the boot media.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv

advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media - ASA A70 and ASA A90

The boot media in your ASA A70 or ASA A90 system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media in the System Management module, and then reinstalling the System Management module.

The boot media is located inside the System Management module and is accessed by removing the module from the system.

Steps

1. Make sure NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged. If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

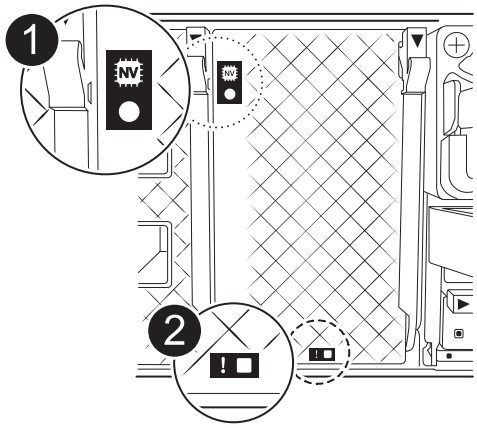

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller's PSUs.

   If your system has DC power, disconnect the power block from the PSUs.

   a. Remove any cables connected to the System Management module. Make sure to label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   c. Depress the system management cam button.
   The cam lever moves away from the chassis.
   d. Rotate the cam lever all the way down and remove the System Management module from the controller module.
   e. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:

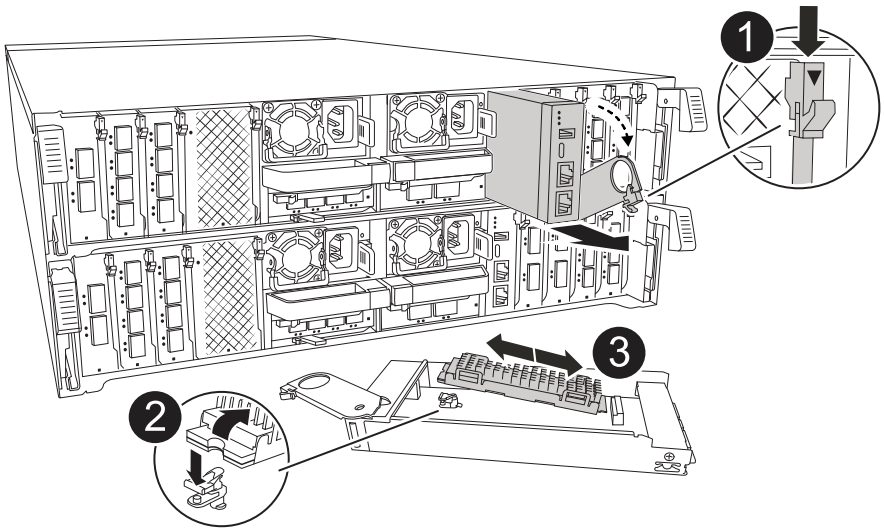

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue locking button.

b. Rotate the boot media up, slide it out of the socket, and set it aside.

5. Install the replacement boot media into the System Management module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down toward the locking button.

c. Push the locking button, rotate the boot media all the way down and then release the locking button.

6. Reinstall the System Management module:

a. Rotate the cable management tray up to the closed position.

b. Recable the System Management module.

7. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

If you have DC power supplies, reconnect the power block to the power supplies.

8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Restore the ONTAP image on the boot media - ASA A70 and ASA A90

After installing the new boot media device in your ASA A70 or ASA A90 system, you can start the automated boot media recovery process to restore the configuration from the

partner node.

During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

   Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

   Show example of configuration error finding prompts

   ```
   Error when fetching key manager config from partner ${partner_ip}:
   ${status}

   Has key manager been configured on this system

   Is the key manager onboard
   ```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed part to NetApp - ASA A70 and ASA A90

If a component in your ASA 70-90 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - ASA A70 and ASA A90

Get started with replacing the chassis of your ASA A70 or ASA A90 storage system by shutting down the controllers, replacing the chassis, and verifying system operations.

1. **Review the chassis replace requirements**

Review the chassis replacement requirements.

2. **Prepare for chassis replace**

Prepare to replace the chassis by locating the system, gathering system credentials and necessary tools, verifying the replacement chassis was received, and labeling the system cables.

3. **Shut down the controllers**

Shut down the controllers so you can perform maintenance on the chassis.

4. **Replace the chassis**

Replace the chassis by moving the components from the impaired chassis to the replacement chassis.

5. **Complete the chassis replacement**

Complete the chassis replacement by bringing the controllers up, giving back the controllers, and returning the failed chassis to NetApp.

Requirement to replace the chassis - ASA A70 and ASA A90

Before replacing the chassis of your ASA A70 or ASA A90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have local

administrator credentials for ONTAP, the correct replacement chassis, and the necessary tools.

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Review the following requirements.

- Make sure all other components in the system are functioning properly; if not, contact NetApp support for assistance.
- Obtain local administrator credentials for ONTAP if you don't have them.
- Make sure that you have the necessary tools and equipment for the replacement.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

What's next?

After reviewing the chassis replacement requirements, you need to prepare for chassis replace.

Prepare to replace chassis - ASA A70 and ASA A90

Prepare to replace the impaired chassis in your ASA A70 or ASA A90 system by identifying the impaired chassis, verifying the replacement components, and labeling the cables and controller modules.

Step 1: Locate and monitor your system

You should open a console session and save sessions logs for future reference, and also turn on the system location LED to find the impaired chassis.

Steps

1. Connect to the serial console port to interface with and monitor the system.
2. Locate and turn on the controller's Location LED:
   a. Use the `system controller location-led show` command to show the current state of the location LED.
   b. Change the state of the location LED to "on":

      `system controller location-led modify -node node1 -state on.`

      The Location LED remains lit for 30 minutes.

Step 2: Verify replacement components

You should verify that you received the necessary components, remove them from packaging, and save the packaging.

Steps

1. Before opening the packaging, you should look at the packaging label and verify:
   - Component part number.
   - Part description.
   - Quantity in the box.
2. Remove the contents from the packaging and use the packaging to returning the failed component to NetApp.

Step 3: Label the cables and controller modules

You should label the cables and controller modules before removing them from the controller modules or chassis.

Steps

1. Label all the cables associated with the storage system. This aids recabling later in this procedure.
2. Label the controller modules.
3. If you are not already properly grounded, ground yourself.

What's next?

After you've prepared to replace your ASA A70 or ASA A90 chassis hardware, you need to shut down the controllers.

Shut down the controllers to replace the chassis - ASA A70 and ASA A90

Shut down the controllers in your ASA A70 or ASA A90 storage system to prevent data loss and ensure system stability when replacing the chassis.

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.

3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

What's next?

After the controllers are shut down, you need to replace the chassis.

Move and replace hardware - ASA A70 and ASA A90

Replace the chassis of your ASA A70 or ASA A90 system when a hardware failure requires it. The replacement process involves removing the controllers and power supply units (PSUs), removing the drives, installing the replacement chassis, and reinstalling the chassis components.

Step 1: Remove the PSUs and cables

You need to remove all four power supply units (PSUs), two per controller, before removing the controller. Removing them lightens the overall weight of each controller.

Steps

1. Remove the four PSUs:
   a. If you are not already grounded, properly ground yourself.
   b. Unplug power cords from the controller module PSU.

      If your system has DC power, disconnect the power block from the PSUs.

   c. Remove the PSU from the controller by rotating the PSU handle up so that you can pull the PSU out, press the PSU locking tab, and then pull PSU out of the controller module.

      The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

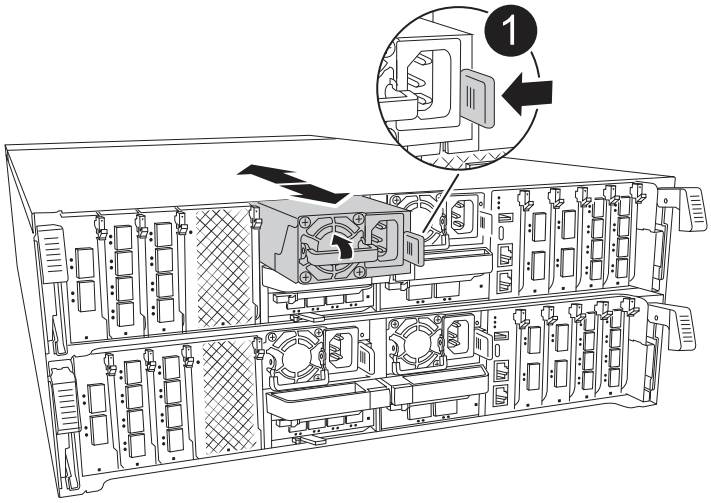

1	Terracotta PSU locking tab

   d. Repeat these steps for the remaining PSUs.
2. Remove the cables:
   a. Unplug the system cables and any SFP and QSFP modules (if needed) from the controller module, but leave them in the cable management device to keep them organized.

      Cables should have been labeled at the beginning of this procedure.

   b. Remove the cable management device from the controller modules and set them aside.

Step 2: Remove the controller modules and drives

Remove the controllers from the chassis and then remove the drives from the chassis.

Steps

1. Check the that amber NVRAM status LED located in slot 4/5 on the back of each controller module is off. Look for the NV icon.

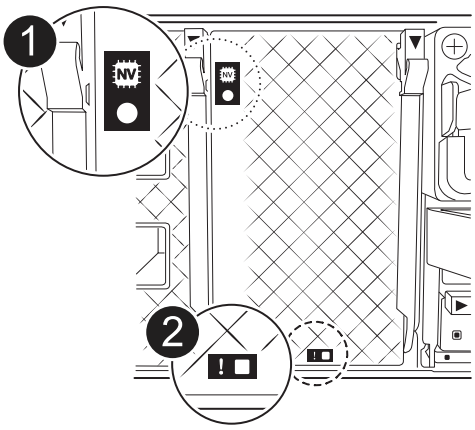

1	NVRAM status LED
2	NVRAM attention LED

- If the NVRAM LED is off, go to the next step.
- If the NVRAM LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact NetApp Support Site for assistance.

2. Remove the controller modules:
   a. Press down on both of the locking latches on the controller, and then rotate both latches downward at the same time.

      The controller module moves slightly out of the chassis.

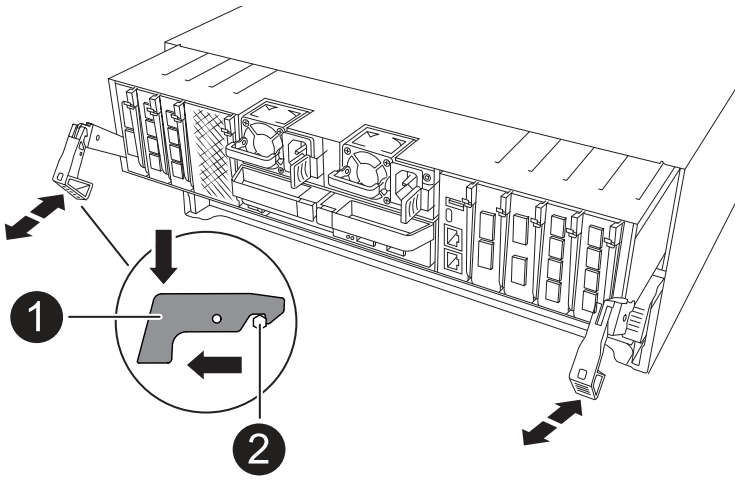

1	Locking latch
2	Locking pin

   b. Slide the controller module out of the chassis by the locking latches, and place it on a flat, stable surface.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

c. Repeat these steps for the second controller module.

3. Remove the drives:
   a. Gently remove the bezel from the front of the system.
   b. Press the release button at the top of the drive carrier face below the LEDs.
   c. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      Drives are fragile. Always use two hands to support the drive weight when removing a drive to prevent damage to them.

   d. Keep track of which drive bay each drive was in and set the drive aside on a static-free cart or table.
   e. Repeat this step for the remaining drives in the chassis.

Step 3: Replace the impaired chassis

Remove the impaired chassis and install the replacement chassis.

Steps

1. Remove the impaired chassis:
   a. Remove the screws from the chassis mount points.
   b. Using two people or a lift, slide the impaired chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
2. Install the replacement chassis:
   a. Using two people or a lift, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
   b. Slide the chassis all the way into the equipment rack or system cabinet.
   c. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.

Step 4: Install the chassis components

After the replacement chassis is installed, you need to install the controller modules, recable them, and then reinstall the drives and PSUs.

Steps

1. Beginning with the bottom controller module, install the controller modules in the replacement chassis:
   a. Align the end of the controller module with the opening in the chassis, and then gently push the controller all the way into the chassis.
   b. Rotate the locking latches upward into the locked position.
   c. If you have not already done so, reinstall the cable management device and recable the controller.

      If you removed the media converters (QSFPs or SFPs), remember to reinstall them.

Make sure that the cables are connected referencing the cable labels.

2. Reinstall the drives into their corresponding drive bays in the front of the chassis.
3. Install all four of the PSUs:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

4. Reconnect the PSU power cables to all four of the PSUs.
   a. Secure the power cable to the PSU using the power cable retainer.

      If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis and secure the power cable to the PSU with the thumbscrews.

The controller modules begin to boot as soon as PSUs are installed and power is restored.

What's next?

After you've replaced the impaired ASA A70 or ASA A90 chassis and reinstalled the components into it, you need to complete the chassis replacement.

Complete the chassis replacement - ASA A70 and ASA A90

Reboot the controllers, verify system health, and return the failed part to NetApp to complete the final step in the ASA A70 and ASA A90 chassis replacement procedure.

Step 1: Boot the controllers and give back the controllers

After the controllers reboot, boot ONTAP and give back the controllers.

Steps

1. Check the console output:
   a. If the controller boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.
   b. If the console displays `waiting for giveback` after the reboot, log into the partner controller and check that the replaced controller is ready for giveback with the `storage failover show` command.
2. Perform the giveback:
   a. Connect the console cable to the partner controller.
   b. Give back the controller with the `storage failover giveback -fromnode local` command.

Step 2: Verify storage system health

After the controller has given back the storage, you should check the overall health with Active IQ Config Advisor.

Steps

1. After the giveback is complete, run Active IQ Config Advisor to verify the health of the storage system.
2. Correct any issues you encounter.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - ASA A70 and ASA A90

Get started with replacing the controller in your ASA A70 or ASA A90 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

Review controller replacement requirements

To replace the controller module, you must meet certain requirements.

Shut down the impaired controller

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

Replace the controller

Replacing the controller includes removing the impaired controller, moving the FRU components to the replacement controller module, and then installing the replacement controller module in the enclosure.

Restore and verify the system configuration

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Recable and give back the controller

Recable the controller and transfer the ownership of storage resources back to the replacement controller.

Complete controller replacement

Verify the LIfs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - ASA A70 and ASA A90

Before replacing the controller in your ASA A70 or ASA A90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements for replacing the controller module.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- Do not use this procedure for controller upgrades; instead, refer to the Choose your controller hardware upgrade procedure for guidance.
- If your system is in a MetroCluster configuration, you must review Choosing the correct recovery procedure to determine whether you should use this procedure.
- You must replace the failed component with the field-replaceable unit (FRU) you received from NetApp.
- You must replace a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text log file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - ASA A70 and ASA A90

Shut down the controller in your ASA A70 or ASA A90 storage system to prevent data loss and ensure system stability when replacing the controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller - ASA A70 and ASA A90

Replace the controller in your ASA A70 or ASA A90 system when a hardware failure requires it. This process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting the system.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

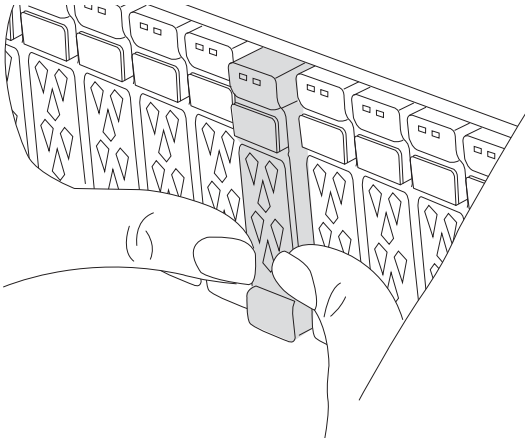

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

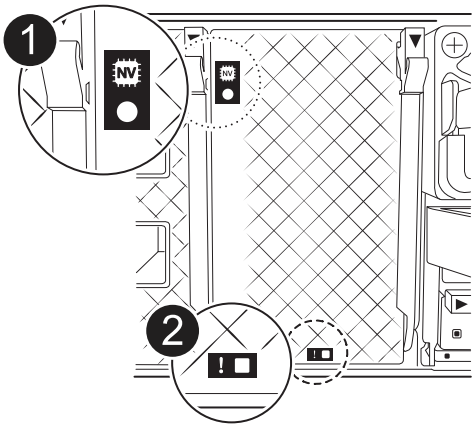

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.

7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

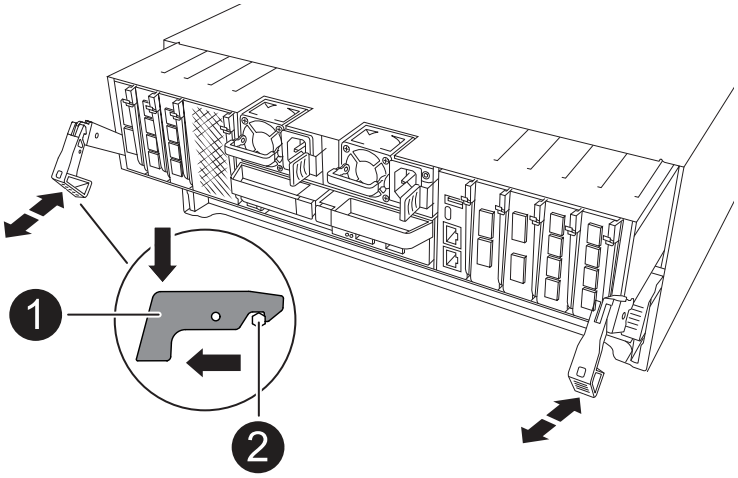

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 2: Move the power supplies

Move the power supplies to the replacement controller.

Steps

1. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

   The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

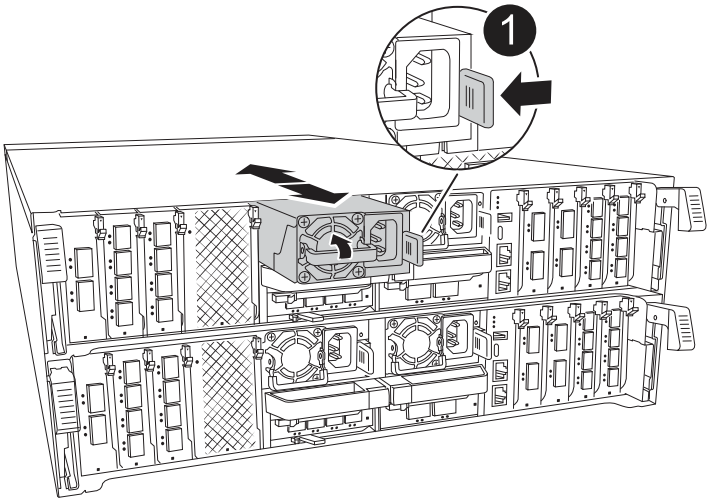

1	Terracotta PSU locking tab
2	Power supply

2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

Move the fans modules to the replacement controller module.

Steps

1. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

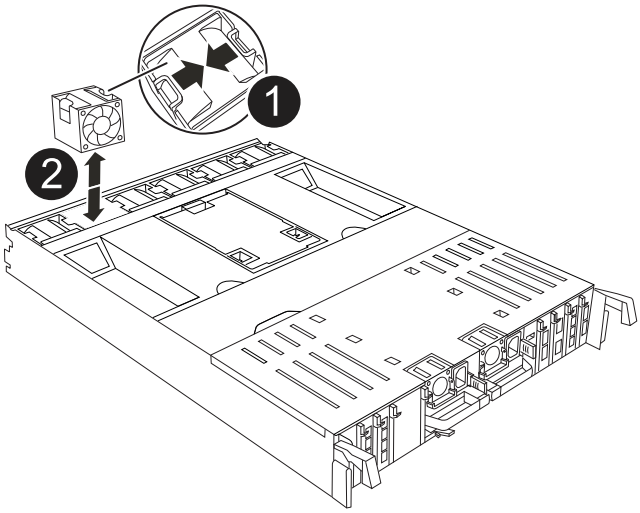

1	Fan locking tabs
2	Fan module

2. Move the fan module to the replacement controller module, and then install the fan module by aligning its edges with the opening in the controller module, and then sliding the fan module into the controller module until the locking latches click into place.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the NV battery

Move the NV battery to the replacement controller module.

Steps

1. Open the air duct cover in the middle of the controller module and locate the NV battery.

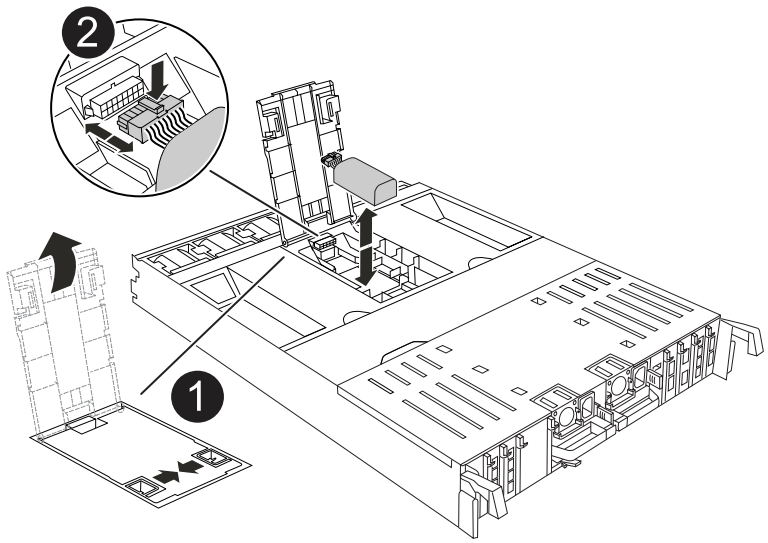

1	NV battery air duct

2	NV battery pack plug

Attention: The NV module LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module.
5. Move the battery pack to the replacement controller module and then install it in the replacement controller module:
   a. Open the NV battery air duct in the replacement controller module.
   b. Plug the battery plug into the socket and make sure that the plug locks into place.
   c. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   d. Close the NV battery air duct.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller module.

Steps

1. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the system DIMMs on the motherboard, using the DIMM map on top of the air duct.

   The DIMM locations, by model, are listed in the following table:

Model	DIMM slot location
FAS70	3, 10, 19, 26
FAS90	3, 7, 10, 14, 19, 23, 26, 30

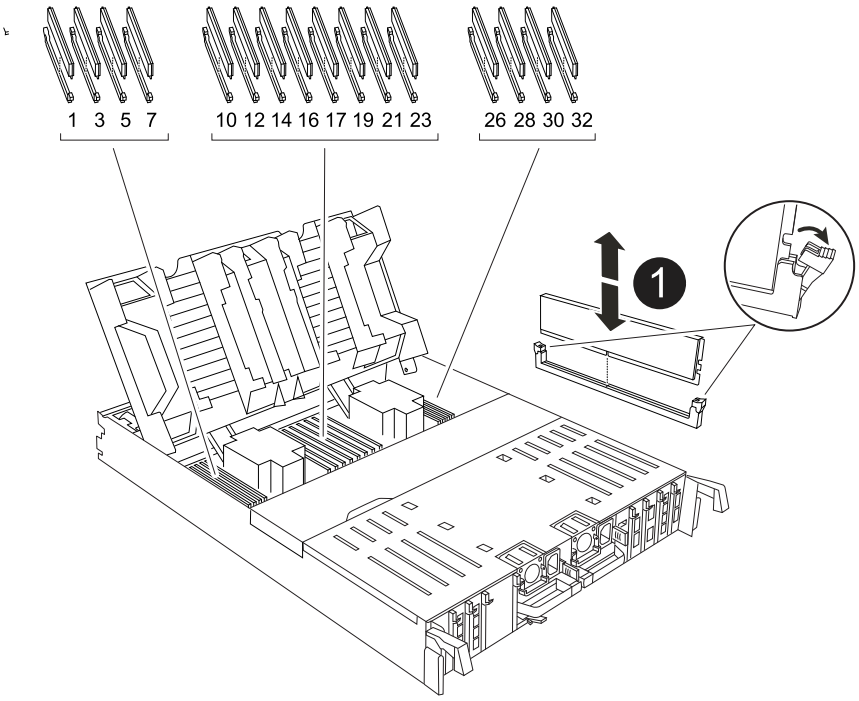

1	System DIMM

3. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Locate the slot on the replacement controller module where you are installing the DIMM.
6. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Repeat these steps for the remaining DIMMs.
9. Close the controller air duct.

Step 6: Move the I/O modules

Move the I/O modules to the replacement controller module.

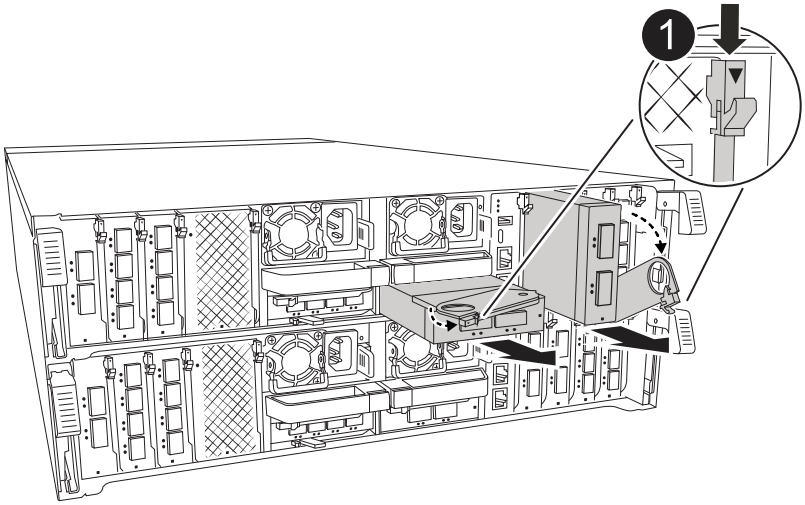

1	I/O module cam lever

Steps

1. Unplug any cabling on the target I/O module.

   Make sure that you label the cables so that you know where they came from.

2. Rotate the cable management arm down by pulling the buttons on the inside of the cable management arm and rotating it down.
3. Remove the I/O modules from the controller module:
   a. Depress the target I/O module cam latch button.
   b. Rotate the cam latch down as far as it will go. For horizontal modules, rotate the cam away from the module as far as it will go.
   c. Remove the module from the controller module by hooking your finger into the cam lever opening and pulling the module out of the controller module.

      Make sure that you keep track of which slot the I/O module was in.

   d. Install the replacement I/O module into the replacement controller module by gently sliding the I/O module into the slot until the I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.
4. Repeat these steps to move the remaining I/O modules, except the modules in slots 6 and 7, to the replacement controller module.

   To move the I/O modules from slots 6 and 7, you must move the carrier containing these I/O modules from the impaired controller module to the replacement controller module.

5. Move the carrier containing the I/O modules in slots 6 and 7 to the replacement controller module:
   a. Press the button on the right-most handle on the carrier handle.
   ..Slide the carrier out of the impaired controller module insert it into the replacement controller module in the same position it was in the impaired controller module.
   b. Gently push the carrier all the way into the replacement controller module until it locks into place.

Step 7: Move the System Management module

Move the System Management module to the replacement controller module.

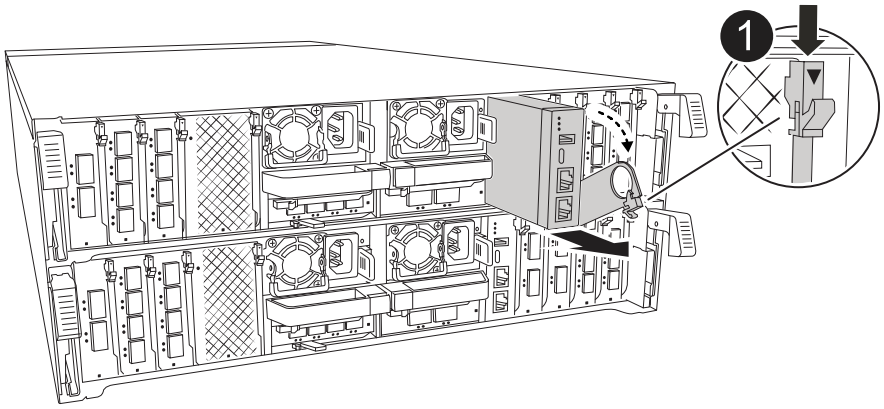

1	System Management module cam latch

Steps

1. Remove the System Management module from the impaired controller module:
   a. Depress the system management cam button.
   b. Rotate the cam lever all the way down.
   c. Loop your finger into the cam lever and pull the module straight out of the system.
2. Install the system management module into the replacement controller module in the same slot that it was in on the impaired controller module:
   a. Align the edges of the System Management module with the system opening and gently push it into the controller module.
   b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

Step 8: Move the NVRAM module

Move the NVRAM module to the replacement controller module.

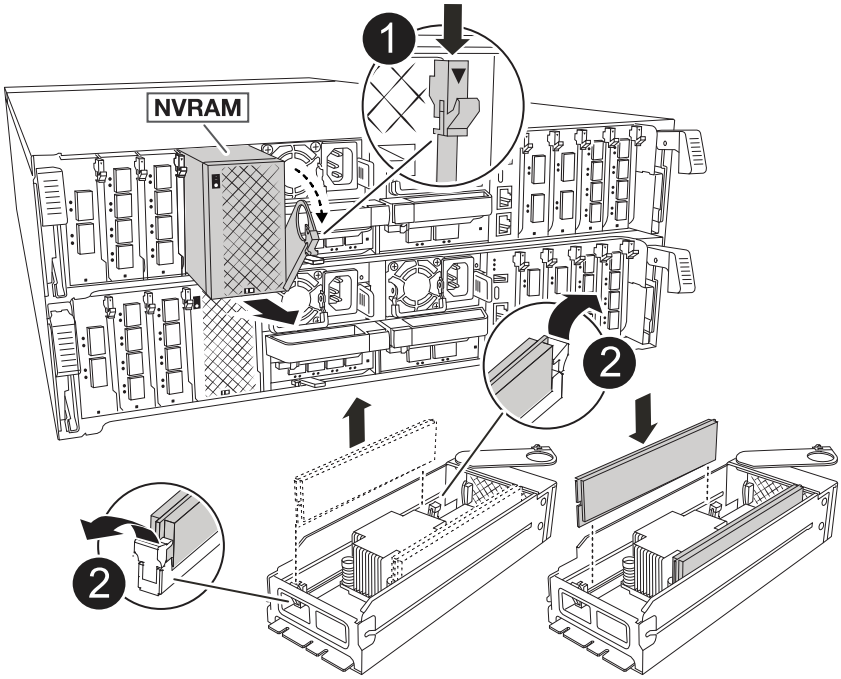

1	Cam locking button
2	DIMM locking tab

Steps

1. Remove the NVRAM module from the impaired controller module:
   a. Depress the cam latch button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.
   c. Remove the NVRAM module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
2. Install the NVRAM module into slot 4/5 in the replacement controller module:
   a. Align the module with the edges of the chassis opening in slot 4/5.
   b. Gently slide the module into the slot all the way, and then push the cam latch all the way up to lock the module in place.

Step 9: Install the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller

module halfway into the system.

> Do not completely insert the controller module in the chassis until instructed to do so.

3. Reinstall the cable management arm, if removed, but do not reconnect any cables to the replacement controller.
4. Plug the console cable into the console port of the replacement controller module and reconnect it to the laptop so that so that it receives console messages when it reboots.
5. Complete the reinstallation of the controller module:

   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

   > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   a. Rotate the locking latches upward into the locked position.
   b. Plug in the power supplies. The controller boots to the LOADER prompt as soon as power is restored.
6. From the LOADER prompt, enter `show date` to display the date and time on the replacement controller. Date and time are in GMT.

   > Time displayed is local time not always GMT and is displayed in 24hr mode.

7. Set the current time in GMT with the `set time hh:mm:ss` command. You can get the current GMT from the partner node the `date -u`command.
8. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

9. Plug the power cords into the power supplies.

   > If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

Restore and verify the system configuration - ASA A70 and ASA A90

Verify that the controller's HA configuration is active and functioning correctly in your ASA A70 or ASA A90 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

Steps

1. Boot to maintenance mode: `boot_ontap maint`

a. Enter `y` when you see *Continue with boot?*.

   If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

5. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha`

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed: `ha-config show`

Step 2: Verify disk list

You must verify the adapter list and paths to all your system disks.

Steps

1. Verify that the adapter lists the paths to all disks with the `storage show disk -p`.

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode: `halt`.

Give back the controller - ASA A70 and ASA A90

Return control of storage resources to the replacement controller so your ASA A70 or ASA A90 system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`
   - If you encounter errors, contact NetApp Support.
9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

    c. Verify status of the keys: `security key-manager key query -restored false`

The output should show no results when when properly synchronized.

If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Complete controller replacement - ASA A70 and ASA A90

To complete the controller replacement for your ASA A70 or ASA A90 system, first restore the NetApp Storage Encryption configuration (if necessary). Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, return the failed part to NetApp.

Step 1: Verify LIFs and check cluster health

Before returning the replacement node to service, verify that the logical interfaces are on their home ports, check the cluster health, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any logical interfaces are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - ASA A70 and ASA A90

Replace a DIMM in your ASA A70 or ASA A90 system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- All other components in the system must be functioning properly; if not, you must contact technical support.
- You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

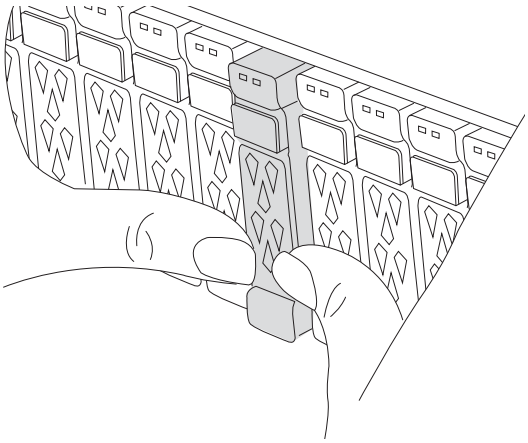

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

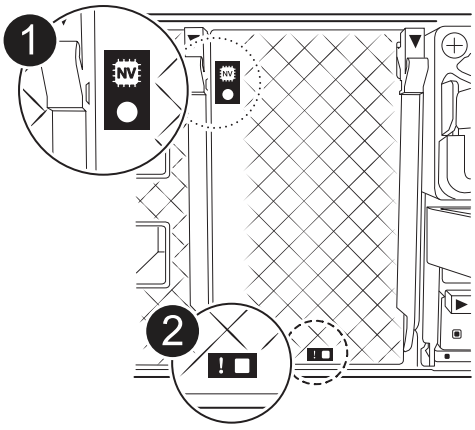

1	NVRAM status LED

2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

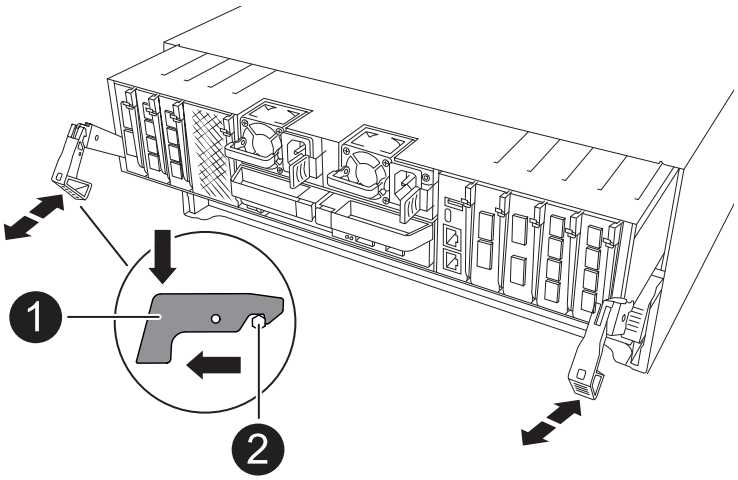

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace a DIMM

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
3. Locate the DIMMs on your controller module and identify the target DIMM.

   Use the FRU map on the controller airduct to locate the DIMM slot.

4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

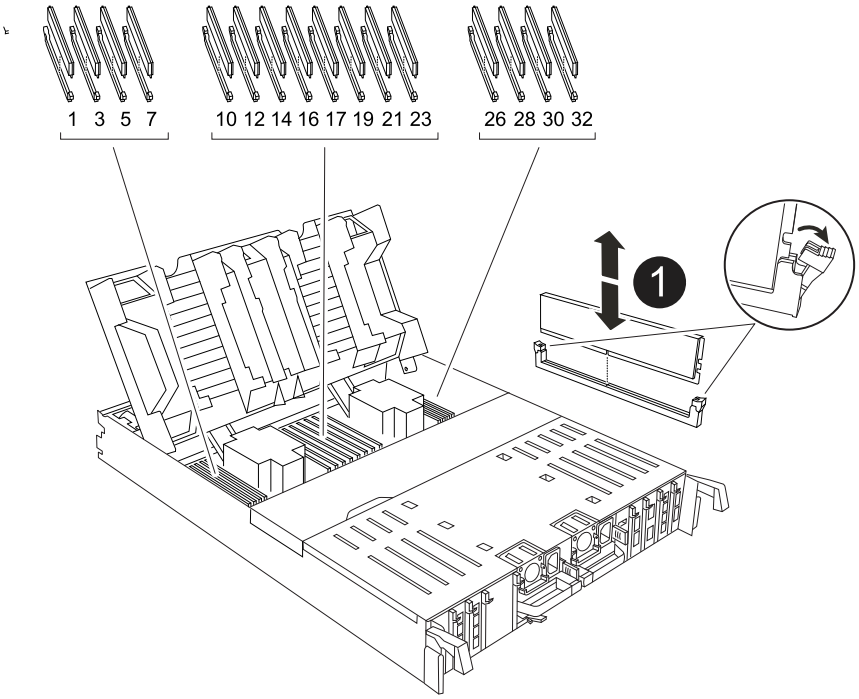

1	DIMM and DIMM ejector tabs

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the controller air duct.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

   Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward into the locked position.
5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

6. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode` *`impaired_node_name`*.

7. If automatic giveback was disabled, reenable it:

   `storage failover modify -node local -auto-giveback true`.

8. If AutoSupport is enabled, restore/unsuppress automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an SSD drive - ASA A70 and ASA A90

Replace a drive in your ASA A70 or ASA A90 system when a drive fails or requires an upgrade. This process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed SSD drive nondisruptively while I/O is in progress.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed disk drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  It can take up to several hours for the drive to appear in the list of failed drives.
- Determine whether SED authentication is enabled.

  How you replace the disk depends on how the disk drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These instructions describe additional steps you must perform before and after replacing an SED.
- Make sure the replacement drive is supported by your platform. See NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

When replacing several disk drives, you must wait 70 seconds between the removal of each failed disk drive and the insertion of the replacement disk drive to allow the storage system to recognize the existence of each new disk.

You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

Steps

1. Properly ground yourself.
2. Remove the bezel from the front of the storage system.
3. Physically identify the failed drive.

- When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the shelf operator display panel and the failed drive illuminate.
- The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

4. Remove the failed drive:

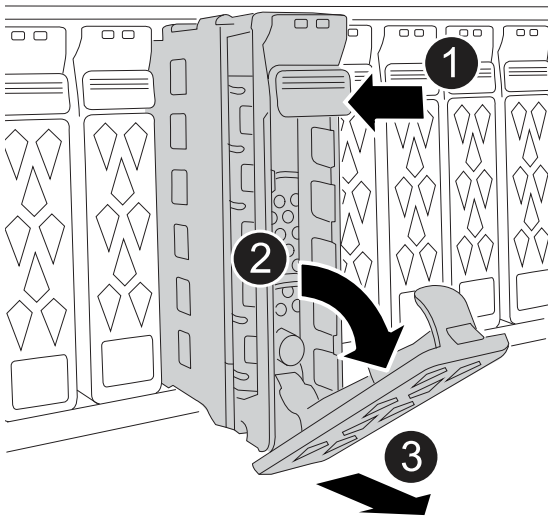

1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Gently push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

7. Verify that the drive's activity (green) LED is illuminated.

   When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is

blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat Steps 3 through 7.
9. Reinstall the bezel on the front of the storage system.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support NetApp Support if you need the RMA number or additional help with the replacement procedure.

Replace a fan module - ASA A70 and ASA A90

Replace a fan module in your ASA A70 or ASA A90 system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

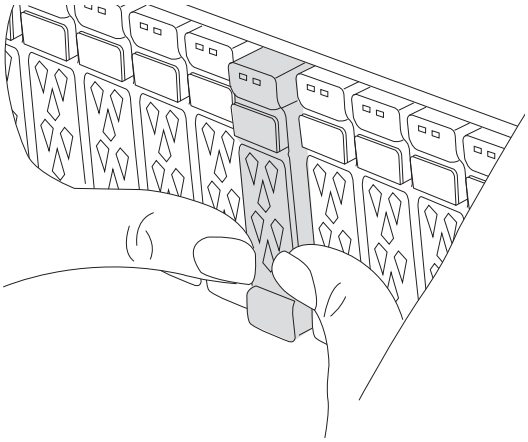

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

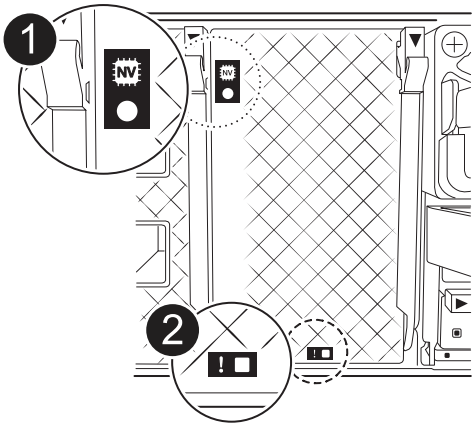

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

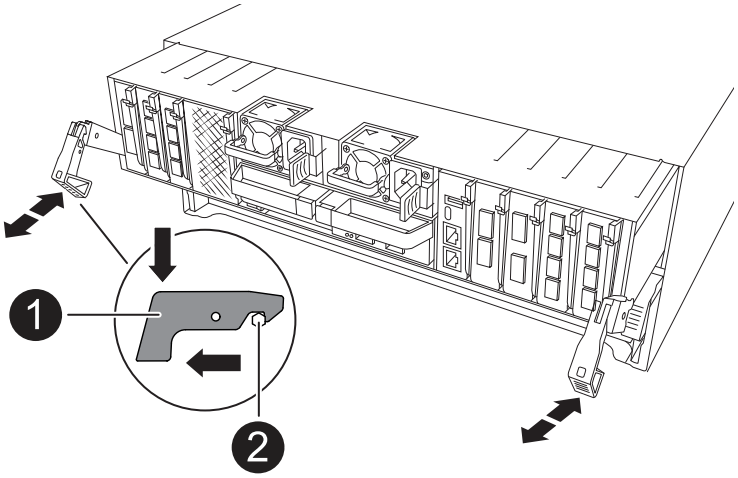

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

Steps

1. Identify the fan module that you must replace by checking the console error messages.
2. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

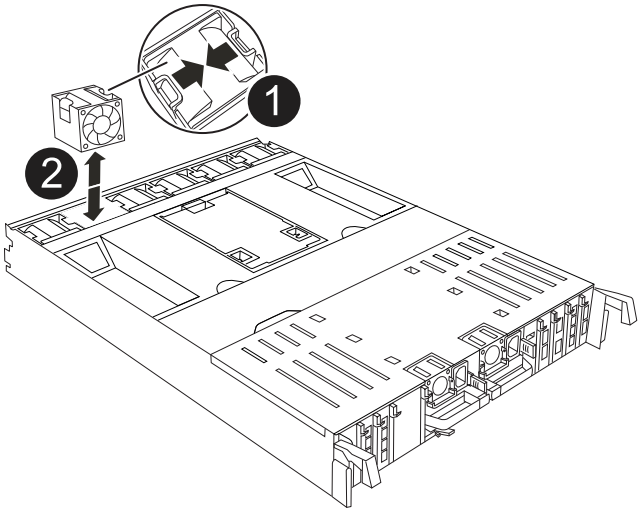

1	Fan locking tabs

2	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module until the locking latches click into place.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

   Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward into the locked position.
5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

6. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode` *`impaired_node_name`*.

7. If automatic giveback was disabled, reenable it:

   `storage failover modify -node local -auto-giveback true`.

8. If AutoSupport is enabled, restore/unsuppress automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END.
```

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace NVRAM - ASA A70 and ASA A90

Replace the NVRAM in your ASA A70 or ASA A90 system when the non-volatile memory becomes faulty or requires an upgrade. The replacement process involves shutting down the impaired controller, replacing the NVRAM module or the NVRAM DIMM, reassigning the disks, and returning the failed part to NetApp.

The NVRAM module consists of the NVRAM12 hardware and field-replaceable DIMMs.
You can replace a failed NVRAM module or the DIMMs inside the NVRAM module.

Before you begin

- Make sure you have the replacement part available. You must replace the failed component with a replacement component you received from NetApp.
- Make sure all other components in the storage system are functioning properly; if not, contact NetApp Support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the NVRAM module or NVRAM DIMM

Replace the NVRAM module or NVRAM DIMMs using the appropriate option.

Option 1: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 4/5 in the chassis and follow the specific sequence of steps.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

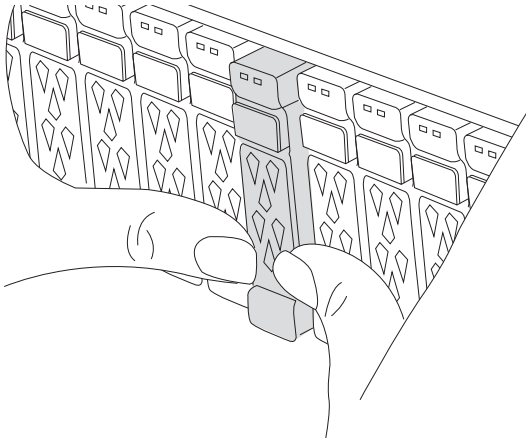

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. If you are not already grounded, properly ground yourself.
4. Unplug the controller's PSUs.

   If your system has DC power, disconnect the power block from the PSUs.

5. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
6. Remove the target NVRAM module from the chassis:
   a. Depress the cam latch button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.
   c. Remove the impaired NVRAM module from the chassis by hooking your finger into the cam lever opening and pulling the module out of the chassis.

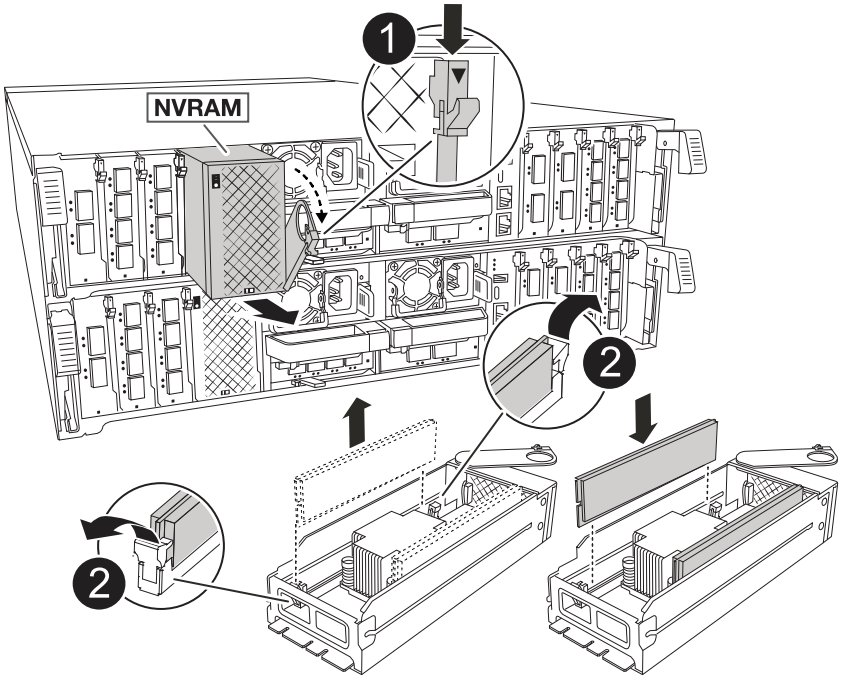

1	Cam locking button
2	DIMM locking tabs

7. Set the NVRAM module on a stable surface.
8. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
9. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 4/5.
   b. Gently slide the module into the slot all the way, and then push the cam latch all the way up to lock the module in place.
10. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

    If you have DC power supplies, reconnect the power block to the power supplies.

11. Rotate the cable management tray up to the closed position.

Option 2: Replace the NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, you must remove the NVRAM module, and then replace the target DIMM.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

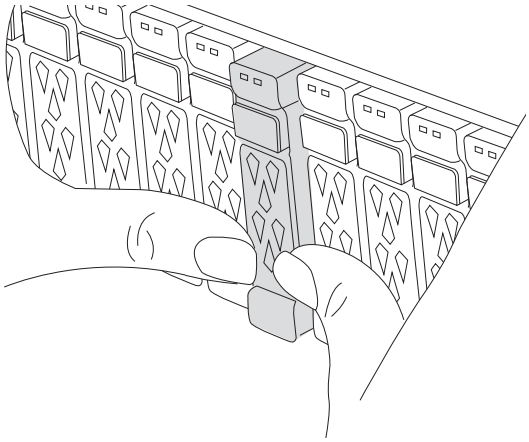

2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller's PSUs.

   If your system has DC power, disconnect the power block from the PSUs.

4. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
5. Remove the target NVRAM module from the chassis:
   a. Depress the cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.
   c. Remove the NVRAM module from the chassis by hooking your finger into the cam lever opening and pulling the module out of the chassis.

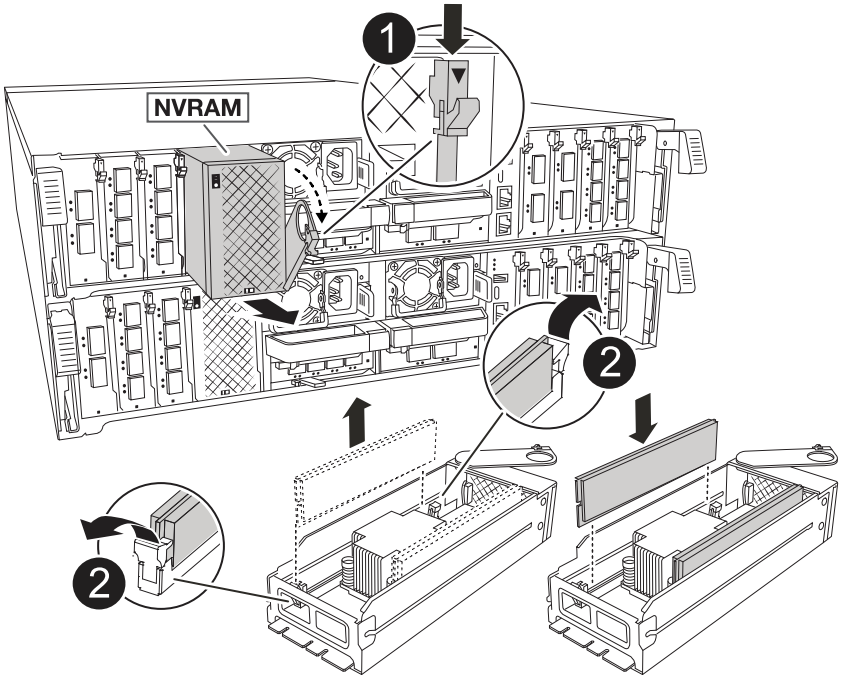

1	Cam locking button
2	DIMM locking tabs

6. Set the NVRAM module on a stable surface.
7. Locate the DIMM to be replaced inside the NVRAM module.

   Consult the FRU map label on the side of the NVRAM module to determine the locations of DIMM slots 1 and 2.

8. Remove the DIMM by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
9. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
10. Install the NVRAM module into the chassis:
    a. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.
11. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

    If you have DC power supplies, reconnect the power block to the power supplies.

12. Rotate the cable management tray up to the closed position.

Step 3: Verify controller state

You must confirm the controller state of the controllers connected to the disk pool when you boot the controller.

Steps

1. If the controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: *halt*
2. From the LOADER prompt on the controller, boot the controller and enter *y* when prompted to override the system ID due to a system ID mismatch.
3. Wait until the Waiting for giveback… message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify the system state: *storage failover show*

   In the command output, you should see a message indicates the state of the controllers.

```
                                Takeover
Node           Partner          Possible State Description
-------------- -------------- --------
-------------------------------------
<nodename>
               <nodename>-    true     Connected to <nodename>-P2-3-178.
               P2-3-178                 Waiting for cluster applications
to
                                        come online on the local node.
AFF-A90-NBC-P2-3-178
               <nodename>-    true     Connected to <nodename>-P2-3-177,
               P2-3-177                 Partial giveback
2 entries were displayed.
```

4. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: *storage failover giveback -ofnode replacement_node_name*

      The controller connects back its storage pool and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter *y*.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: *storage failover show*
5. Verify all disks are displayed: `storage disk show`

```
::> storage disk show
                     Usable           Disk    Container   Container
Disk                   Size Shelf Bay Type    Type        Name
---------------- ---------- ----- --- ------- ----------- ---------
1.0.0                3.49TB     0   0 SSD-NVM aggregate   pod_NVME_SSD_1
1.0.1                3.49TB     0   1 SSD-NVM aggregate   pod_NVME_SSD_1
1.0.2                3.49TB     0   2 SSD-NVM aggregate   pod_NVME_SSD_1
1.0.3                3.49TB     0   3 SSD-NVM aggregate   pod_NVME_SSD_1
1.0.4                3.49TB     0   4 SSD-NVM aggregate   pod_NVME_SSD_1

[...]
48 entries were displayed.
```

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NV battery - ASA A70 and ASA A90

Replace the NV battery in your ASA A70 or ASA A90 system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.
-
To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

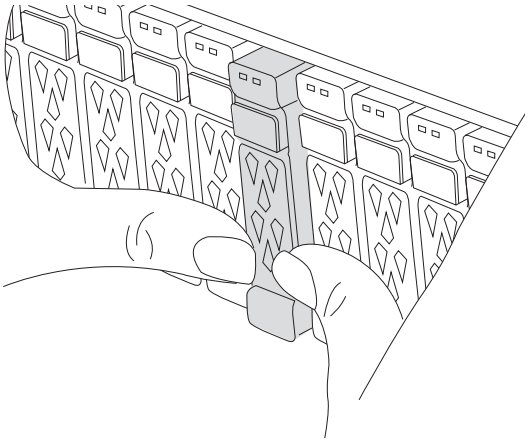

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

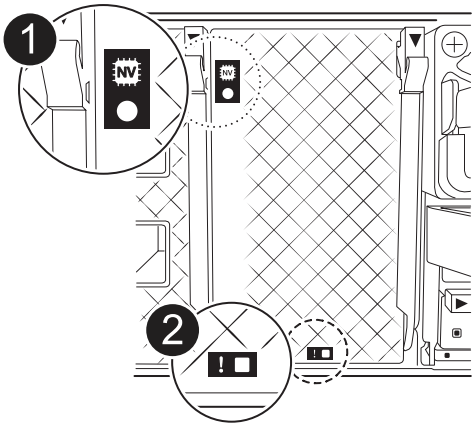

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

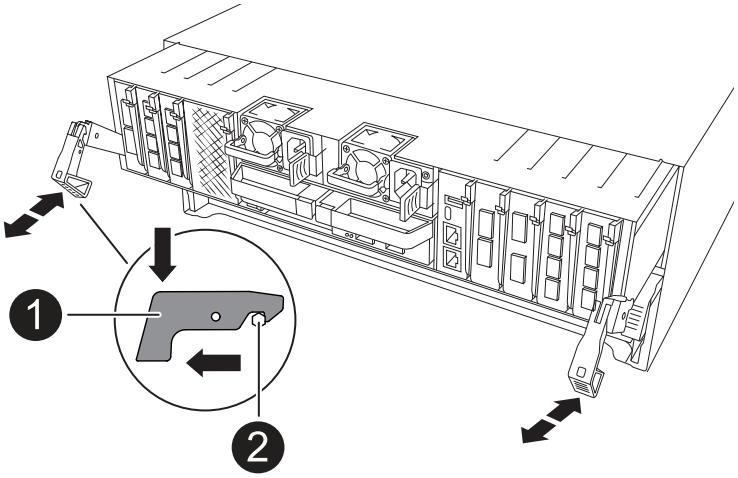

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the NV battery

Remove the failed NV battery from the controller module and install the replacement NV battery.

Steps

1. Open the air duct cover and locate the NV battery.

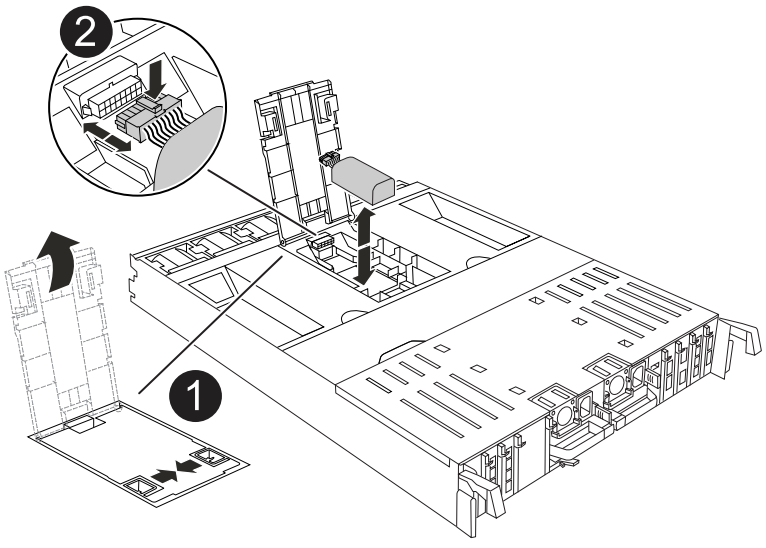

1	NV battery air duct cover

2	NV battery plug

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module, and then set it aside.
5. Remove the replacement battery from its package.
6. Install the replacement battery pack into the controller:
   a. Plug the battery plug into the riser socket and make sure that the plug locks into place.
   b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
7. Close the NV air duct cover.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   ⓘ Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

   Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward into the locked position.
5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

6. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name.`

7. If automatic giveback was disabled, reenable it:

   `storage failover modify -node local -auto-giveback true.`

8. If AutoSupport is enabled, restore/unsuppress automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END.`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Overview of add and replace an I/O module - ASA A70 and ASA A90

The ASA A70 and ASA A90 systems offer flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your ASA A70 or ASA A90 storage system with the same type of I/O module, or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

- Add an I/O module

  Adding additional modules can improve redundancy, helping to ensure that the system remains operational even if one module fails.

- Replace an I/O module

  Replacing a failing I/O module can restore the system to its optimal operating state.

I/O slot numbering

The I/O slots on ASA A70 and ASA A90 controllers are numbered 1 through 11, as shown in the following illustration.

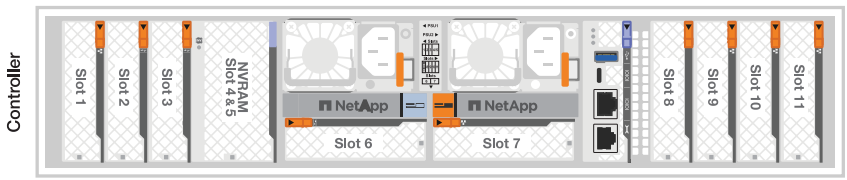

Add an I/O module to your ASA A70 or ASA A90 system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your ASA A70 and ASA A90 storage system when there are empty slots available or when all slots are fully populated.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has two location LEDs, one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- Make sure that all other components are functioning properly.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the target slot blanking module from the carrier:
   a. Depress the cam latch on the blanking module in the target slot.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
4. Install the I/O module:
   a. Align the I/O module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
5. Cable the I/O module to the designated device.

   Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

6. Rotate the cable management tray up to the closed position.
7. From the LOADER prompt, reboot the node:

   ```
   bye
   ```

   This reinitializes the I/O module and other components and reboots the node.

8. Give back the controller from the partner controller:

   ```
   storage failover giveback -ofnode target_node_name
   ```

9. Repeat these steps for controller B.
10. From the healthy node, restore automatic giveback if you disabled it:

    ```
    storage failover modify -node local -auto-giveback true
    ```

11. If AutoSupport is enabled, restore automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the target I/O module from the chassis:
   a. Depress the cam latch button.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

      Make sure that you keep track of which slot the I/O module was in.
5. Install the I/O module into the target slot in the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
6. Cable the I/O module to the designated device.
7. Repeat the remove and install steps to replace additional modules for the controller.
8. Rotate the cable management tray up to the closed position.
9. Reboot the controller from the LOADER prompt:_bye_

   This reinitializes the PCIe cards and other components and reboots the node.

   If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement
10. Give back the controller from the partner controller:

    ```
    storage failover giveback -ofnode target_node_name
    ```
11. Enable automatic giveback if it was disabled:

```
storage failover modify -node local -auto-giveback true
```

12. Do one of the following:
    - If you removed a NIC I/O module and installed a new NIC I/O module, use the following network command for each port:

      ```
      storage port modify -node *<node name> -port *<port name> -mode network
      ```

    - If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
13. Repeat these steps for controller B.

Replace an I/O module - ASA A70 and ASA A90

Replace an I/O module in your ASA A70 or ASA A90 system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

ou can use this procedure with all versions of ONTAP supported by your storage system.

Before you begin

- You must have the replacement part available.
- Make sure all other components in the storage system are functioning properly; if not, contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace a failed I/O module

Steps

To replace an I/O module, locate it within the controller module and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.

   Make sure to label the cables so that you know where they came from.

3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the I/O module from the controller module:

   This following illustration shows removing a horizontal and vertical I/O module. Typically, you will only remove one I/O module.

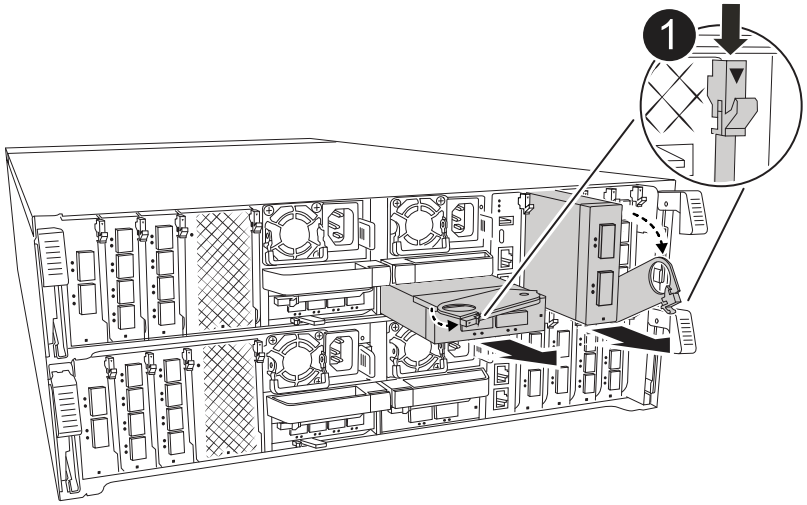

1	Cam locking button

a. Depress the cam latch button.

b. Rotate the cam latch do away from the module as far as it will go.

c. Remove the module from the controller module by hooking your finger into the cam lever opening and pulling the module out of the controller module.

   Make sure that you keep track of which slot the I/O module was in.

5. Set the I/O module aside.
6. Install the replacement I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently slide the module into the slot all the way into the controller module, and then rotate the cam latch all the way up to lock the module in place.
7. Cable the I/O module.
8. Repeat the remove and install steps to replace additional modules for the controller.
9. Rotate the cable management tray into the locked position.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```

3. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

4. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - ASA A70 and ASA A90

Replace an AC or DC power supply unit (PSU) in your ASA A70 or ASA A90 system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cable, replacing the faulty PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

About this task

This procedure is written for replacing one PSU at a time.

> Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

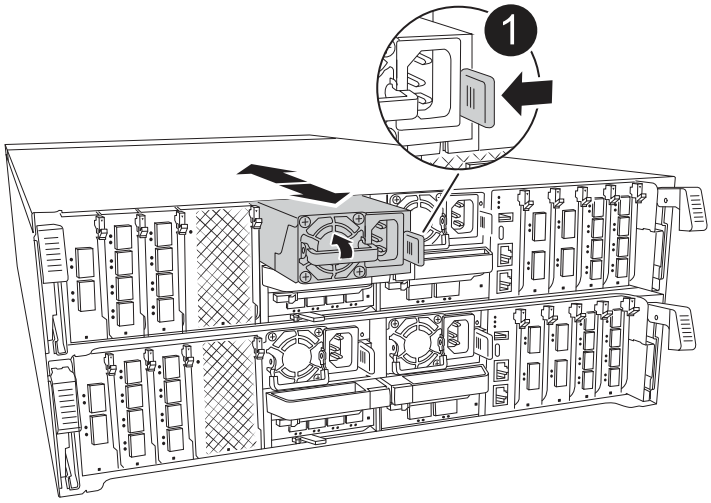

1	Terracotta PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.
6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU.

b. Secure the power cable to the PSU using the power cable retainer.

Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

> The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

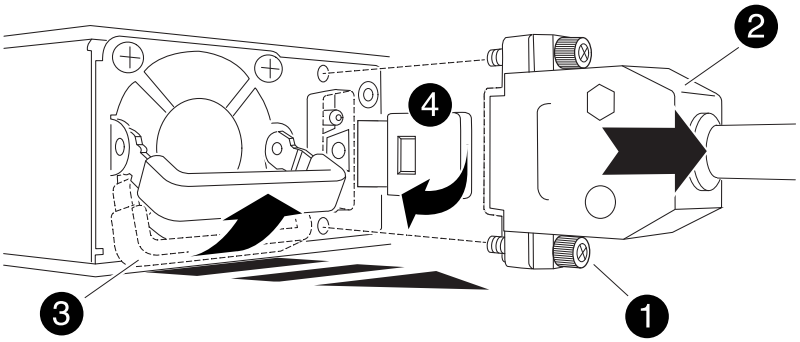

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle
4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.

b. Gently push the PSU into the controller module until the locking tab clicks into place.

The power supplies will only properly engage with the internal connector and lock in place one way.

> To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

   Once power is restored to the PSU, the status LED should be green.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA A70 and ASA A90

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your ASA A70 or ASA A90 system to ensure that services and applications relying on accurate time synchronization remain operational.

You can use this procedure with all versions of ONTAP supported by your system.

Make sure all other components in the system are functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

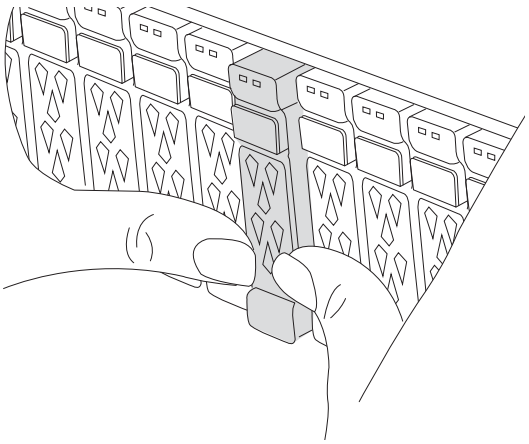

2. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

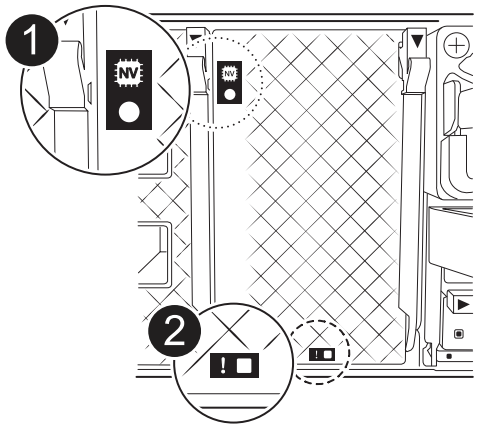

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
3. If you are not already grounded, properly ground yourself.
4. Unplug the controller module power supply cables from the controller module power supplies (PSU).

   If your system has DC power, disconnect the power block from the PSUs.

5. Unplug the system cables and SFP and QSFP modules (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

6. Remove the cable management device from the controller module.
7. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

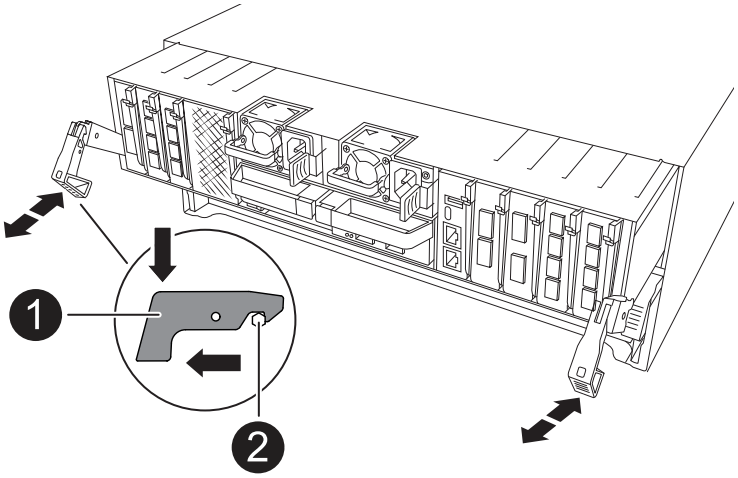

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the RTC battery

Remove failed RTC battery and install the replacement RTC battery.

Steps

1. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the RTC battery under the air duct.

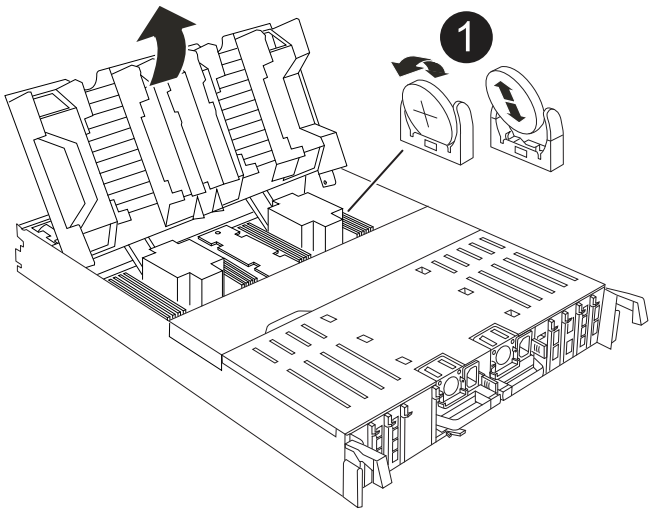

1	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

   Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

   Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward into the locked position.
5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

6. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name.`

7. If automatic giveback was disabled, reenable it:

   `storage failover modify -node local -auto-giveback true.`

8. If AutoSupport is enabled, restore/unsuppress automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END.`

Step 5: Reset the time and date on the controller

After you replace the RTC battery, insert the controller, and power on for the first BIOS reset, you will see the following error messages:

```
RTC date/time error. Reset date/time to default

RTC power failure error
```

These messages are expected and you can continue with this procedure.

Steps

1. Check the date and time on the healthy controller with the `cluster date show` command.

   If your system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*

   a. At the LOADER prompt on the target controller, check the time and date with the `cluster date show` command.

   b. If necessary, modify the date with the `set date mm/dd/yyyy` command.

   c. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.

2. Confirm the date and time on the target controller.
3. At the LOADER prompt, enter *bye* to reinitialize the PCIe cards and other components and let the controller reboot.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the system management module - ASA A70 and ASA A90

Replace the System Management module in your ASA A70 or ASA A90 system when it becomes defective or its firmware is corrupted. The replacement process involves shutting down the controller, replacing the failed System Management module, rebooting the controller, updating the license keys, and returning the failed part to NetApp.

Before you begin

- This procedure uses the following terminology:
  - The impaired controller is the controller on which you are performing maintenance.
  - The healthy controller is the HA partner of the impaired controller.
- All other system components must be working properly.
- The partner controller must be able to take over the impaired controller.
- You must replace the failed component with a replacement FRU component you received from your provider.

About this task

The System Management module, located at the back of the controller in slot 8, contains onboard components for system management, as well as ports for external management. The target controller must be shut down to replace an impaired System Management module or replace the boot media.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the System Management module

Replace the impaired system management module.

Steps

1. Make sure NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged. If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

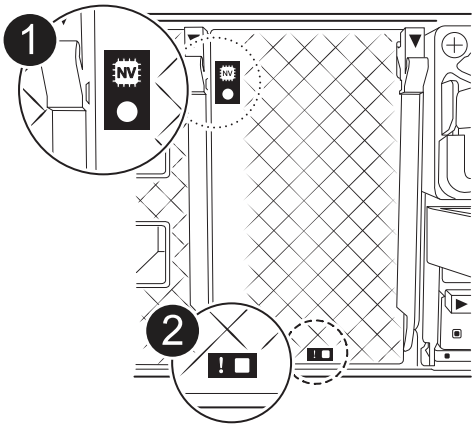

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.
3. Unplug the controller's PSUs.

If your system has DC power, disconnect the power block from the PSUs.

4. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
5. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.

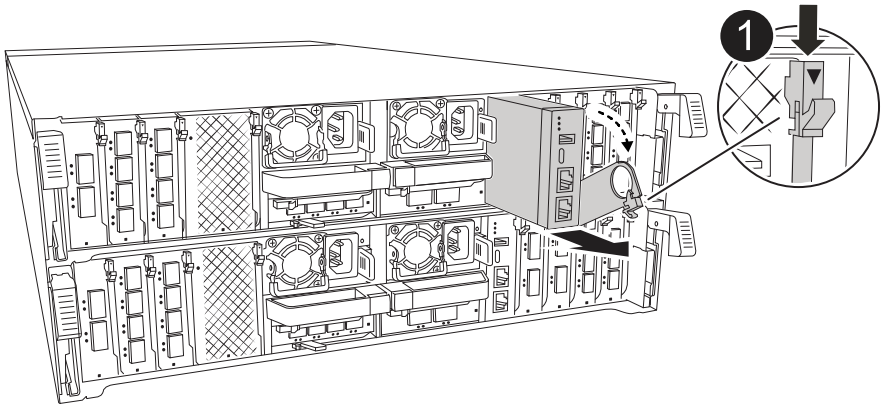

1	System Management module cam latch

6. Remove the System Management module:
   a. Depress the system management cam button.
   The cam lever moves away from the chassis.
   b. Rotate the cam lever all the way down.
   c. Loop your finger into the cam lever and pull the module straight out of the system.
   d. Place the System Management module on an anti-static mat, so that the boot media is accessible.
7. Move the boot media to the replacement System Management module:

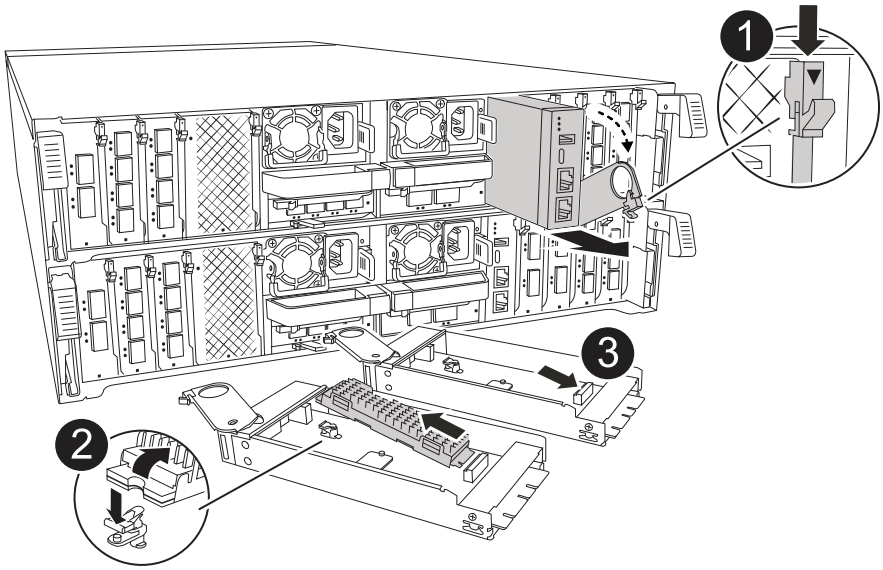

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue locking button.
The boot media rotates slightly upward.

b. Rotate the boot media up, slide it out of the socket.

c. Install the boot media in the replacement System Management module:

i. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

ii. Rotate the boot media down toward until it engages the locking button. Depress the blue locking if necessary.

8. Install the system management module:

a. Align the edges of the replacement System Management module with the system opening and gently push it into the controller module.

b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

9. Recable the System Management module.

10. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

If you have DC power supplies, reconnect the power block to the power supplies.

11. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

Reboot the controller module.

Steps

1. Enter *bye* at the LOADER prompt.
2. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

3. Restore automatic giveback:

```
storage failover modify -node local -auto-giveback true
```

4. If an AutoSupport maintenance window was triggered, end it:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Step 4: Install licenses and register serial number

You must install new licenses for the node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the node. However, if the node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the for the node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`
4. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA A20, ASA A30, and ASA A50 systems

Overview of hardware maintenance - ASA A20, ASA A30, and ASA A50

Maintain the hardware of your ASA A20, ASA A30, or ASA A50 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the ASA A20, ASA A30, and ASA A50 systems have already been deployed as a storage nodes in the ONTAP environment.

System components

For the ASA A20, ASA A30, and ASA A50 storage systems, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller and drives.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
NV battery	The non-volatile memory (NV) battery is responsible for providing power to the NVMEM components while data in-flight is being destaged to flash memory after a power loss.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.

Boot media

Boot media replacement workflow - ASA A20, ASA A30, and ASA A50

Get started with replacing the boot media in your ASA A30, ASA A20, or ASA A50 storage system by reviewing the replacement requirements, shutting down the impaired controller, replacing the boot media, restoring the image on the boot media, and verifying the system functionality.

1. **Review the boot media requirements**

   Review the requirements for boot media replacement.

2. **Shut down the impaired controller**

   Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

3. **Replace the boot media**

   Remove the failed boot media from the impaired controller and install the replacement boot media.

4. **Restore the image on the boot media**

   Restore the ONTAP image from the healthy controller.

5. **Return the failed part to NetApp**

   Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements to replace the boot media - ASA A20, ASA A30, and ASA A50

Before replacing the boot media in your ASA A20, ASA A30 or ASA A50 storage system, ensure you meet the necessary requirements and considerations for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0M (wrench) port on the impaired controller is working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

Review the following requirements.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.

- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the impaired controller.

Shut down the controller to replace the boot media - ASA A20, ASA A30, and ASA A50

Shut down the impaired controller in your ASA A20, ASA A30, or ASA A50 storage system to prevent data loss and ensure system stability when replacing the boot media.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media - ASA A20, ASA A30, and ASA A50

The boot media in your ASA A20, ASA A30, or ASA A50 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then reinstalling the controller module.

About this task

If needed, you can turn on the platform chassis location (blue) LEDs to aid in physically locating the affected platform. Log into the BMC using SSH and enter the `system location-led` *`on`* command.

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led` *`off`* command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.

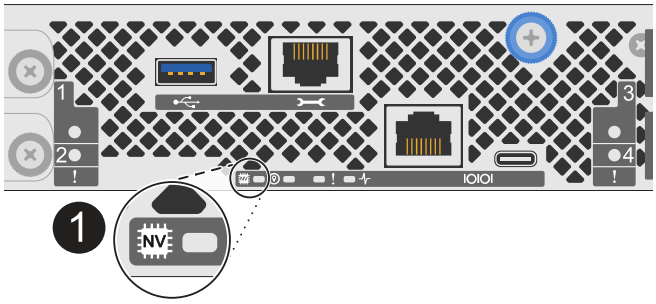

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

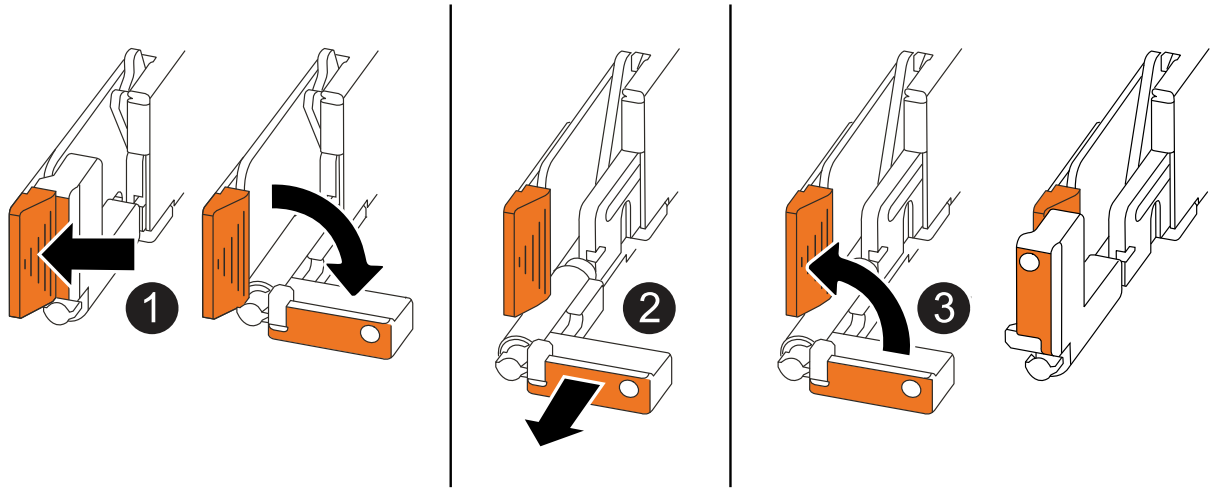

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.
7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:

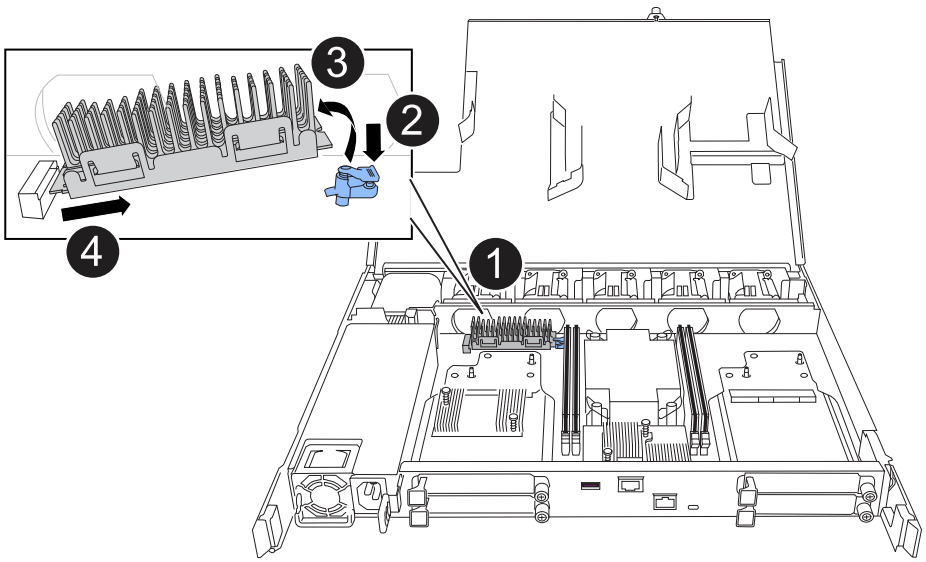

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:
   a. Remove the boot media from its package.
   b. Slide the socket end of the boot media into its socket.
   c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

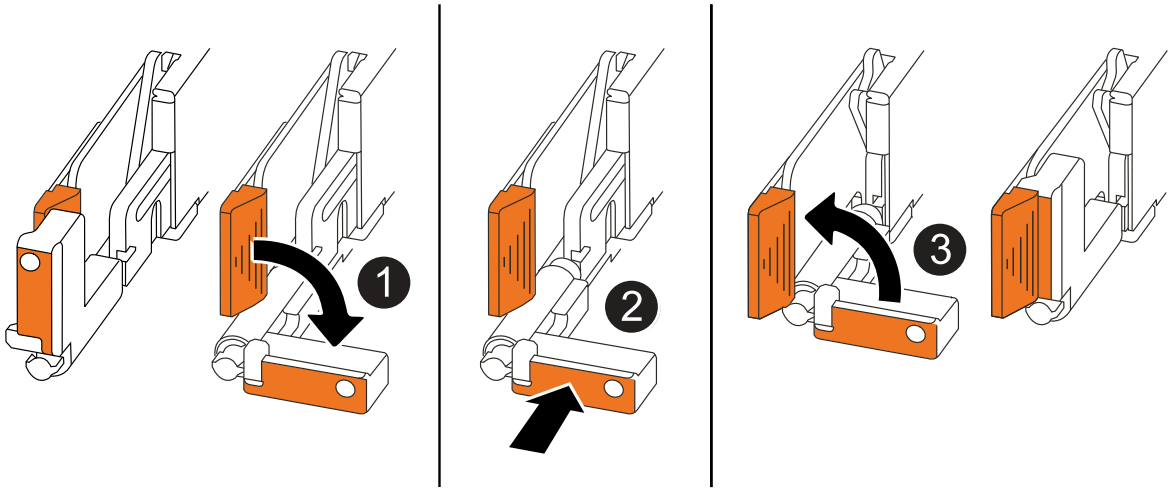

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.

   Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

      The controller boots to the LOADER prompt when fully seated in the chassis. It gets its power from the partner controller.

b. Rotate the controller handles up and lock in place with the tabs.

5. Reconnect the power cord to the PSU on the impaired controller.

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next

After physically replacing the impaired boot media, you restore the ONTAP image from the partner node.

Restore the ONTAP image on the boot media - ASA A20, ASA A30, and ASA A50

After installing the new boot media device in your ASA A20, ASA A30, or ASA A50 storage system, you can start the automated boot media recovery process to restore the configuration from the healthy node.

During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed part to NetApp - ASA A20, ASA A30, and ASA A50

When a component in your ASA A20, ASA A30, or ASA A50 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - ASA A20, ASA A30, and ASA A50

Get started with replacing the chassis of your ASA A20, ASA A30, or ASA A50 storage system by reviewing the replacement requirements, shutting down the controllers, replacing the chassis, and verifying system operations.

Review the chassis replace requirements

Review the requirements to replace the chassis.

Shut down the controllers

Shut down the controllers so you can perform maintenance on the chassis.

Replace the chassis

Replace the chassis by moving the drives and any drive blanks, controllers (with the power supplies), and bezel from the impaired chassis to the new chassis, and swapping out the impaired chassis with the new chassis of the same model as the impaired chassis.

Complete chassis replacement

Verify the HA state of the chassis and return the failed part to NetApp.

Requirements to replace the chassis - ASA A20, ASA A30, and ASA A50

Before replacing the chassis of your ASA A20, ASA A30, or ASA A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement chassis, and the necessary tools.

Review the following requirements and considerations.

Requirements

- The replacement chassis must be the same model as the impaired chassis. This procedure is for a like-for-like replacement, not for an upgrade.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Considerations

- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your storage system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, drives, any drive blanks, and controllers to the new chassis.

What's next?

After you've reviewed the requirements to replace the chassis, you need to shut down the controllers.

Shut down the controllers to replace the chassis - ASA A20, ASA A30, and ASA A50

Shut down the controllers in your ASA A20, ASA A30, or ASA A50 storage system to prevent data loss and ensure system stability when replacing the chassis.

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

```
system node autosupport invoke -node * -type all -message "MAINT=2h Replace
```

```
chassis"
```

5. Identify the SP/BMC address of all cluster nodes:

```
system service-processor show -node * -fields address
```

6. Exit the cluster shell:

```
exit
```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict -sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

```
Warning: Are you sure you want to halt node <node_name>? {y|n}:
```

10. Wait for each controller to halt and display the LOADER prompt.

What's next?

After you've shut down the controllers, you need to replace the chassis.

Replace the chassis - ASA A20, ASA A30, and ASA A50

Replace the chassis of your ASA A20, ASA A30, or ASA A50 storage system when a hardware failure requires it. The replacement process involves removing the controllers, removing the drives, installing the replacement chassis, and reinstalling the chassis components.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

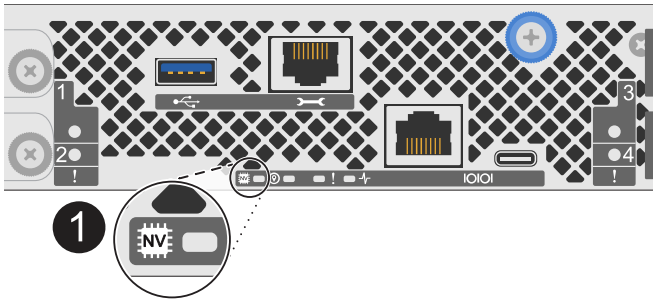

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

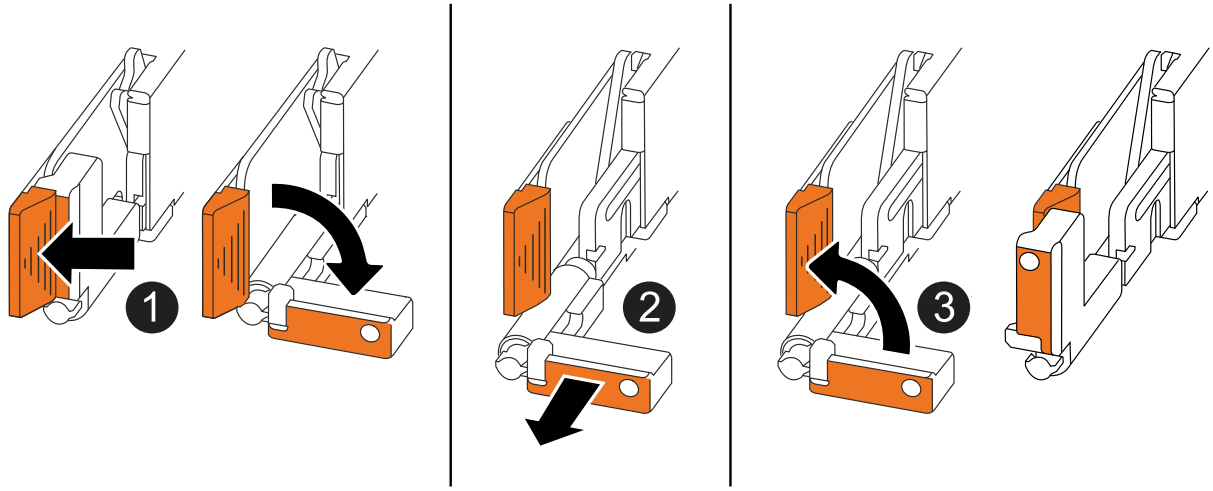

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Repeat these steps for the other controller in the chassis.

Step 2: Remove the drives from the impaired chassis

You need to remove all of the drives and any drive blanks from the impaired chassis so that later in the procedure you can install them in the replacement chassis.

1. Gently remove the bezel from the front of the storage system.
2. Remove the drives and any drive blanks:

   Keep track of what drive bay each drive and drive blank was removed from because they must be installed in the same drive bays in the replacement chassis.

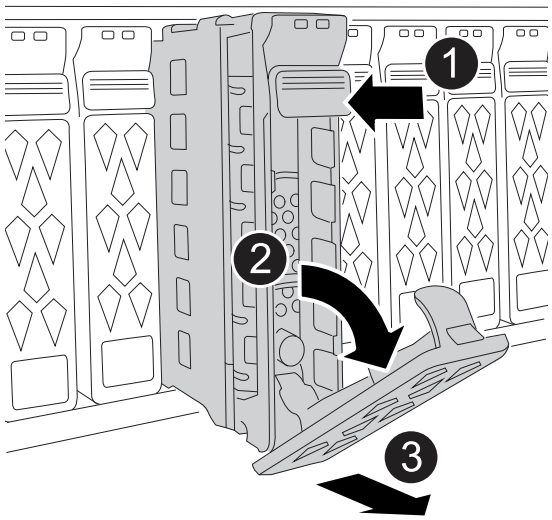

1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

3. Set the drives aside on a static-free cart or table.

Step 2: Replace the chassis from within the equipment rack or system cabinet

You remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis, install the drives, any drive blanks, and then install the bezel.

1. Remove the screws from the impaired chassis mount points.

   Set the screws aside to use later in this procedure.

   If the storage system shipped in a NetApp system cabinet, you must remove additional screws at the rear of the chassis before the chassis can be removed.

2. Using two people or a power lift, remove the impaired chassis from the equipment rack or system cabinet by sliding it off the rails, and then set it aside.
3. Using two people, install the replacement chassis into the equipment rack or system cabinet by sliding it onto the rails.
4. Secure the front of the replacement chassis to the equipment rack or system cabinet using the screws you removed from the impaired chassis.

Step 4: Install the controllers and drives

Install the controllers and drives into the replacement chassis and reboot the controllers.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when installing a controller, and can be used as a reference for the rest of the controller installation steps.

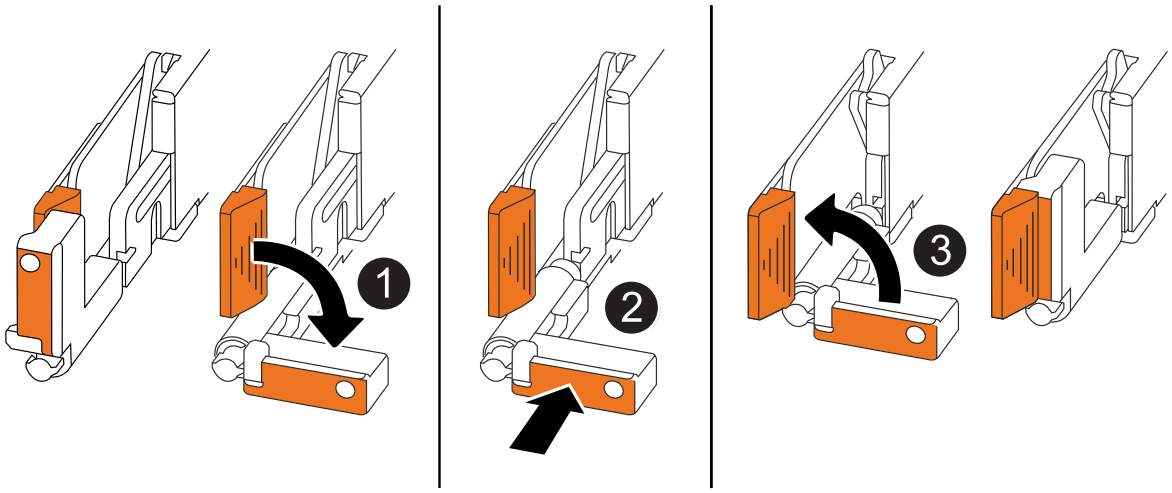

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis and push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

1. Insert one of the controllers into the chassis:
   a. Align the back of the controller with the opening in the chassis.
   b. Firmly push on the handles until the controller meets the midplane and is fully seated in the chassis.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   c. Rotate the controller handles up and lock in place with the tabs.
2. Recable the controller, as needed, except for the power cords.
3. Repeat these steps to install the second controller into the chassis.
4. Install the drives and any drive blanks you removed from the impaired chassis into the replacement chassis:

   The drives and drive blanks must be installed in the same drive bays in the replacement chassis.

   a. With the cam handle in the open position, use both hands to insert the drive.

b. Gently push until the drive stops.

c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

d. Repeat the process for the remaining drives.

5. Install the bezel.
6. Reconnect the power cords to the power supplies (PSU) in the controllers.

   Once power is restored to a PSU, the status LED should be green.

   The controllers begin to boot as soon as the power is restored.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. If controllers boot to the LOADER prompt, reboot the controllers:

   ```
   boot_ontap
   ```

8. Turn AutoSupport back on:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Complete chassis replacement - ASA A20, ASA A30, and ASA A50

Verify the HA state of the chassis and then return the failed part to NetApp to complete the final step in the ASA A20, ASA A30, and ASA A50 chassis replacement procedure.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your storage system configuration.

1. In Maintenance mode, from either controller, display the HA state of the local controller and chassis:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your storage system configuration:
   a. Set the HA state for the chassis:

```
ha-config modify chassis HA-state
```

The value for HA-state should be *ha*.
The value for HA-state can be one of the following:
* ***ha***
* *mcc* (not supported in ASA)

b. Confirm that the setting has changed:

```
ha-config show
```

3. If you have not already done so, recable the rest of your storage system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - ASA A20, ASA A30, and ASA A50

Get started with replacing the controller in your ASA A20, ASA A30, or ASA A50 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

Review the controller replacement requirements

Review the requirements to replace the controller.

Shut down the impaired controller

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

Replace the controller

Replacing the controller includes removing the impaired controller, moving FRU components to the replacement controller, installing the replacement controller in the chassis, setting the time and date, and then recabling.

Restore and verify the system configuration

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Give back the controller

Transfer the ownership of storage resources back to the replacement controller.

Complete controller replacement

Verify the LIFs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - ASA A20, ASA A30, and ASA A50

Before replacing the controller in your ASA A20, ASA A30, or ASA A50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

You must review the requirements and considerations for the controller replacement procedure.

Requirements

- All shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- You must replace a controller with a controller of the same model type. You cannot upgrade your system by just replacing the controller.
- You cannot change any drives or shelves as part of this procedure.
- You must always capture the controller's console output to a text log file.

  The console output provides you with a record of the procedure you can use to troubleshoot issues you might encounter during the replacement process.

Considerations

- It is important that you apply the commands in this procedure to the correct controller:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.

What's next?

After you've reviewed the requirements to replace the impaired controller, you need to shut down the impaired controller.

Shut down the impaired controller - ASA A20, ASA A30, and ASA A50

Shut down the impaired controller in your ASA A20, ASA A30, or ASA A50 storage system to prevent data loss and ensure system stability when replacing the controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After you've shut down the impaired controller, you need to replace the controller.

Replace the controller - ASA A20, ASA A30, and ASA A50

Replace the controller in your ASA A20, ASA A30, or ASA A50 storage system when a hardware failure requires it. The replacement process involves removing the impaired

controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

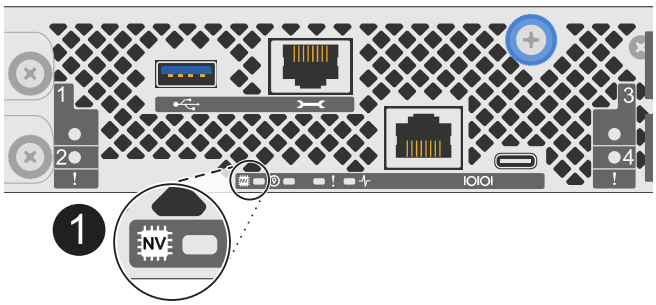

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

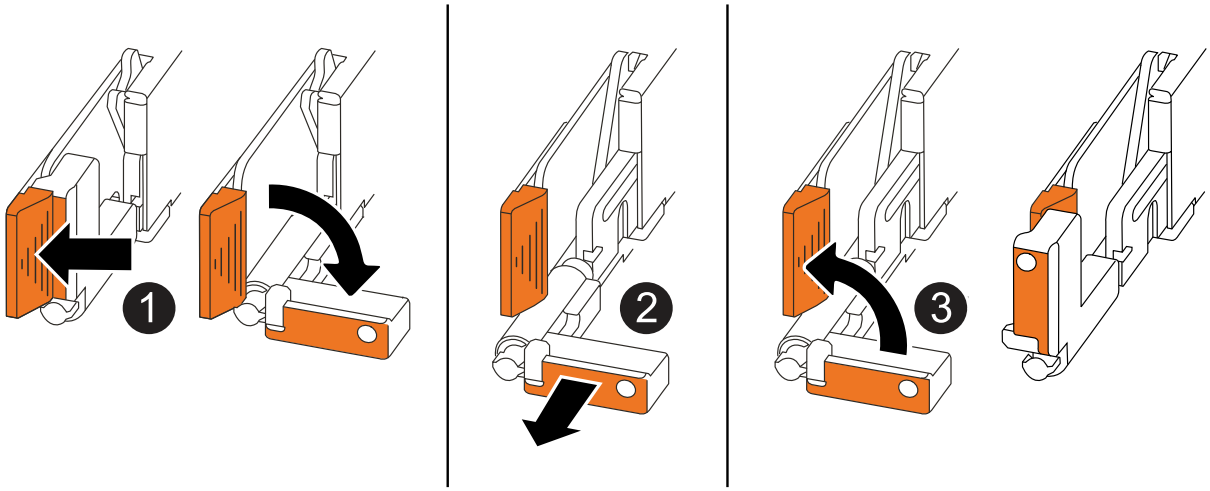

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Move the power supply

Move the power supply (PSU) to the replacement controller.

1. Move the PSU from the impaired controller:

   Make sure the left side controller handle is in the upright position to allow you access to the PSU.

Option 1: Move an AC PSU

To move an AC PSU, complete the following steps.

Steps

1. Remove the AC PSU from the impaired controller:

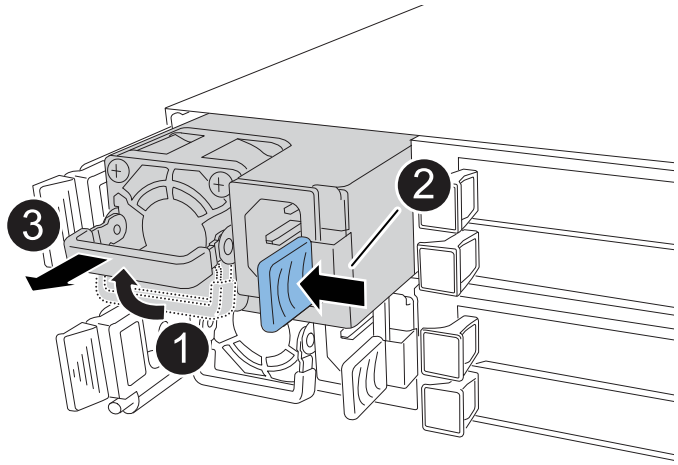

1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

2. Insert the PSU into the replacement controller:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.

Option 2: Move a DC PSU

To move a DC PSU, complete the following steps.

Steps

1. Remove the DC PSU from the impaired controller:

a. Rotate the handle up, to its horizontal position, and then grasp it.
b. With your thumb, press the terracotta tab to release the locking mechanism.
c. Pull the PSU out of the controller while using your other hand to support its weight.

> The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.

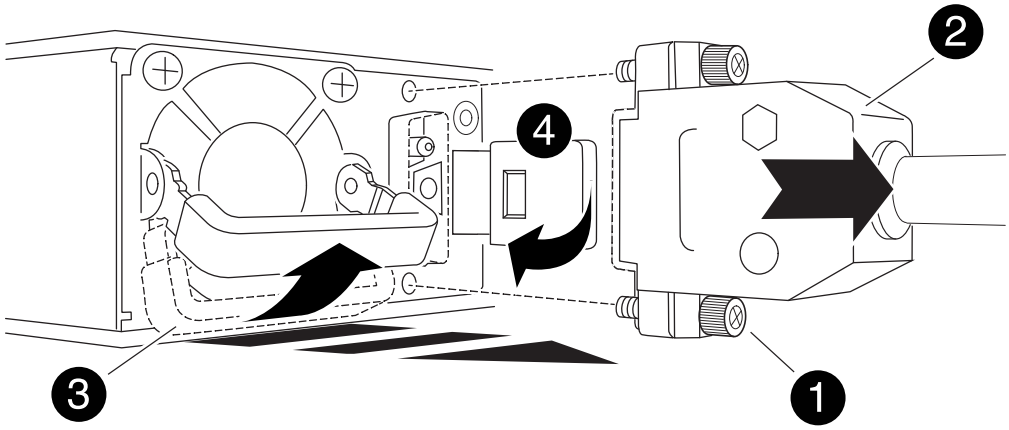

1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

2. Insert the PSU into the replacement controller:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently slide the PSU into the controller until the locking tab clicks into place.

      A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

      > To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.

Step 3: Move the fans

Move the fans to the replacement controller.

1. Remove one of the fans from the impaired controller:

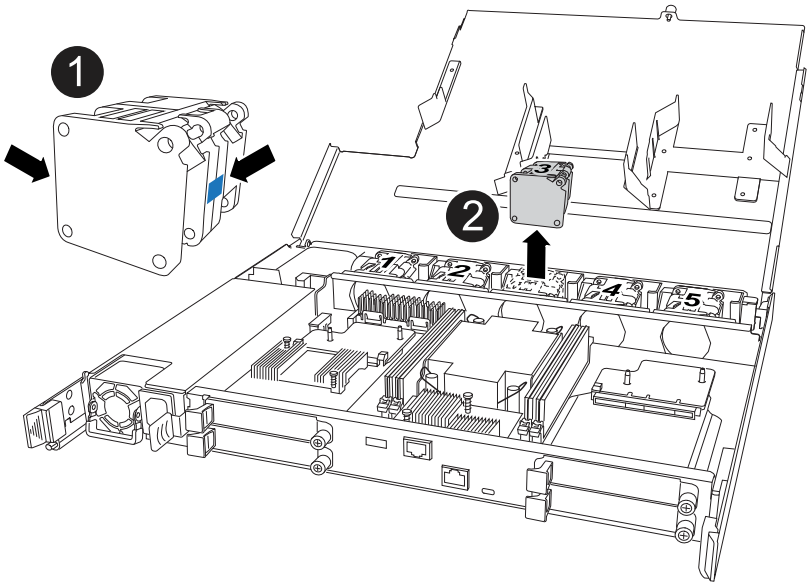

1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

2. Insert the fan into the replacement controller by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.
3. Repeat these steps for the remaining fans.

Step 4: Move the NV battery

Move the NV battery to the replacement controller.

1. Remove the NV battery from the impaired controller:

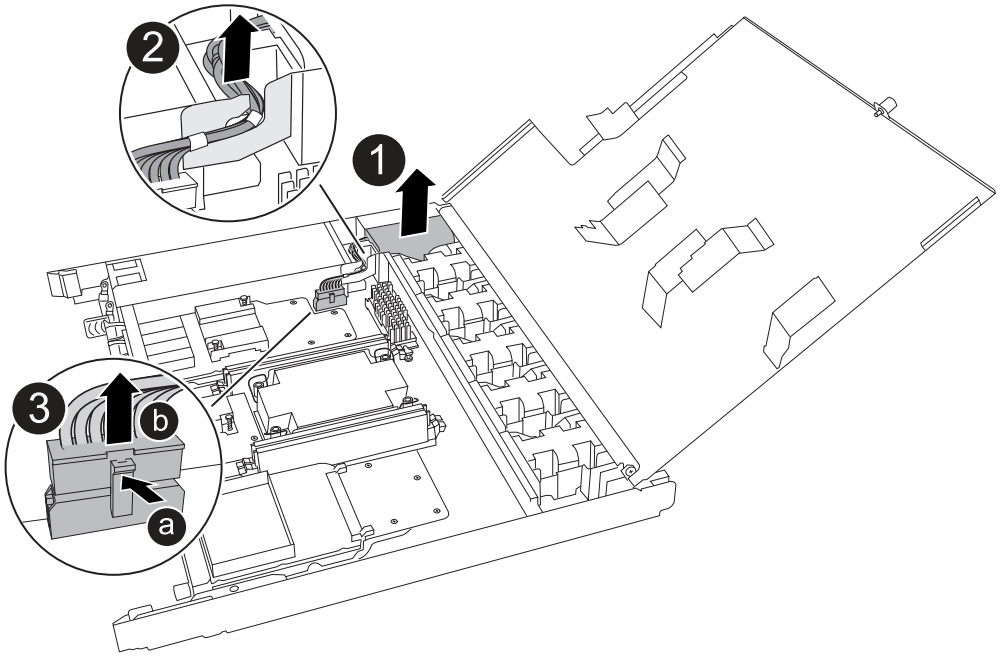

1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

2. Install the NV battery into the replacement controller:
   a. Plug the wiring connector into its socket.
   b. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   c. Place the NV battery into the compartment.

      The NV battery should sit flush in its compartment.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller.

If you have DIMM blanks, you do not need to move them, the replacement controller should come with them installed.

1. Remove one of the DIMMs from the impaired controller:

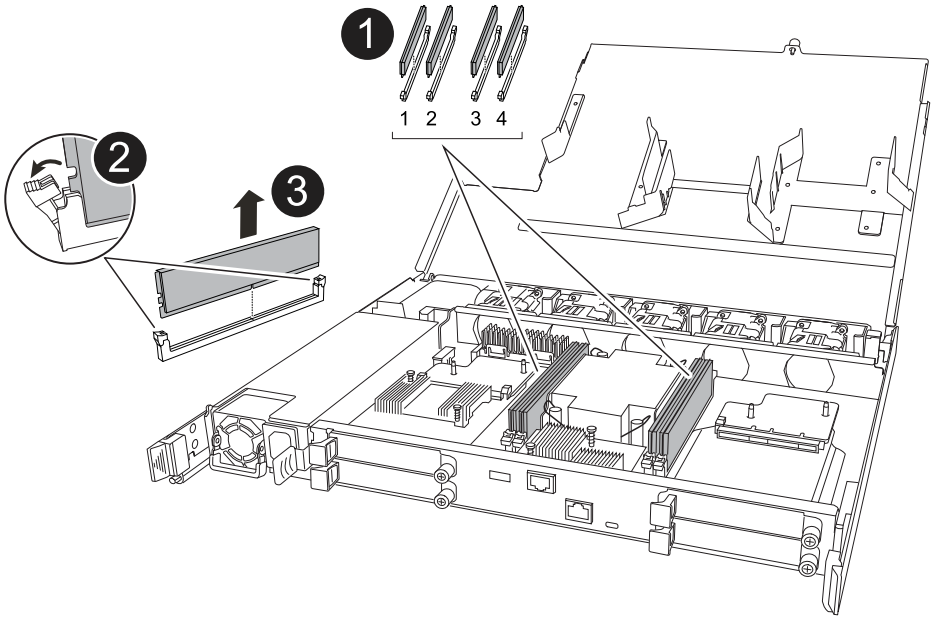

1	DIMM slot numbering and positions. Depending on your storage system model, you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller in the proper orientation. • Eject the DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

2. Install the DIMM in the replacement controller:
   a. Make sure that the DIMM ejector tabs on the connector are in the open position.
   b. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

      When inserted correctly, the DIMM goes in easily but fits tightly in the slot. If not, reinsert the DIMM.

   c. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
   d. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.
3. Repeat these steps for the remaining DIMMs.

Step 6: Move the boot media

Move the boot media to the replacement controller.

1. Remove the boot media from the impaired controller:

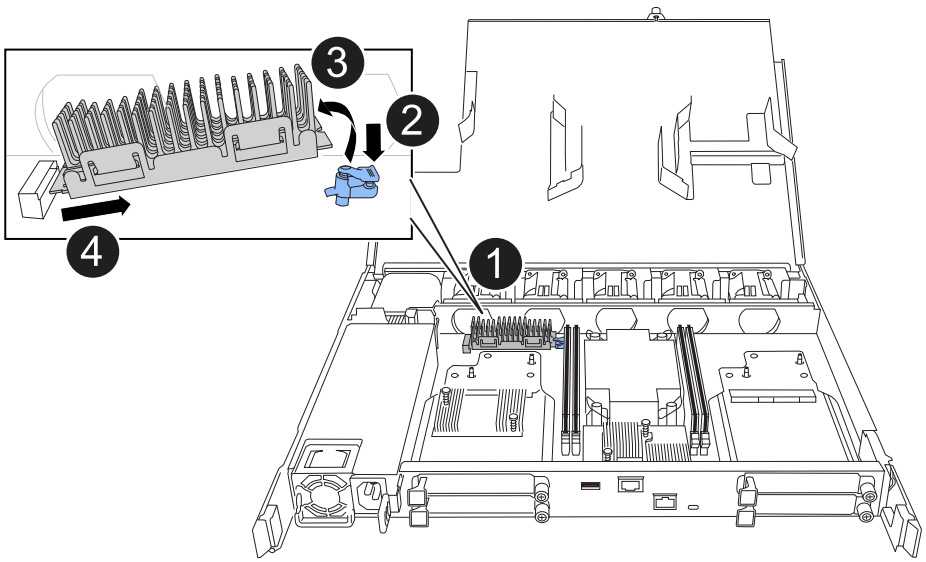

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

2. Install the boot media into the replacement controller:
   a. Slide the socket end of the boot media into its socket.
   b. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 7: Move the I/O modules

Move the I/O modules and any I/O blanking modules to the replacement controller.

1. Unplug cabling from one of the I/O modules.

   Make sure to label the cables so that you know where they came from.

2. Remove the I/O module from the impaired controller:

   Make sure that you keep track of which slot the I/O module was in.

   If you are removing the I/O module in slot 4, make sure the right side controller handle is in the upright position to allow you access to the I/O module.

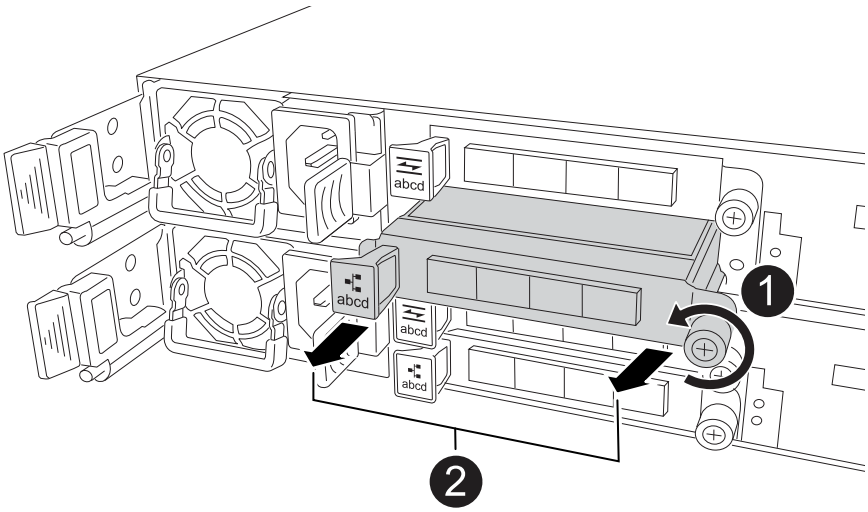

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

3. Install the I/O module into the replacement controller:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
4. Repeat these steps to move the remaining I/O modules and any I/O blanking modules to the replacement controller.

Step 8: Install the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

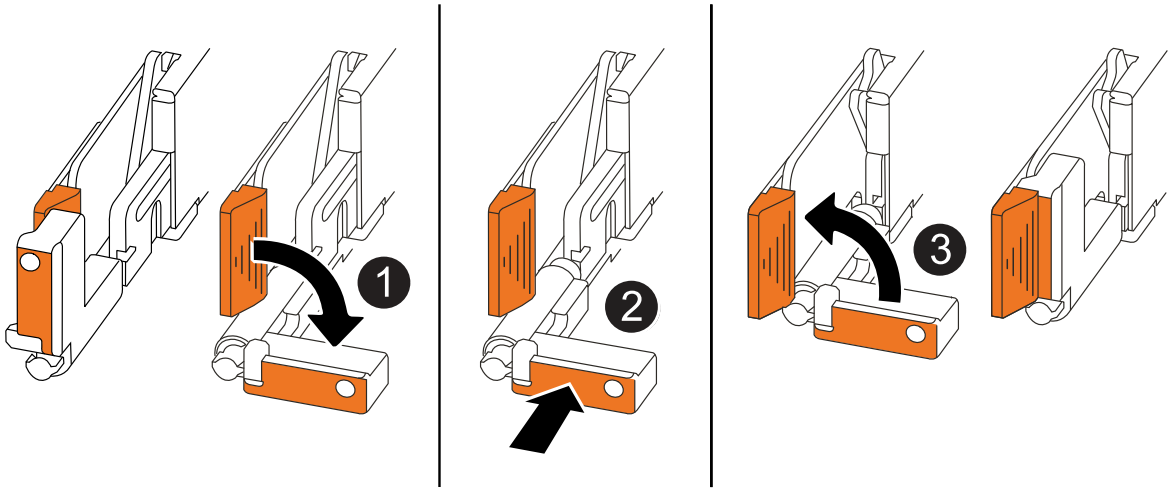

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Take the controller to the LOADER prompt by pressing CTRL-C to abort AUTOBOOT.
6. Set the time and date on the controller:

   Make sure you are at the controller's LOADER prompt.

   a. Display the date and time on the controller:

      ```
      show date
      ```

      Time and date default is in GMT. You have the option to display in local time and in 24hr mode.

   b. Set the current time in GMT:

      ```
      set time hh:mm:ss
      ```

      You can get the current GMT from the healthy node:

      ```
      date -u
      ```

   c. Set the current date in GMT:

      ```
      set date mm/dd/yyyy
      ```

      You can get the current GMT from the healthy node:
      ```
      date -u
      ```
7. Recable the controller as needed.
8. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

After you've replaced the impaired controller, you need to restore the system configuration.

Restore and verify the system configuration - ASA A20, ASA A30, and ASA A50

Verify that the controller's HA configuration is active and functioning correctly in your ASA A20, ASA A30, or ASA A50 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

You must verify the `HA` state of the controller and, if necessary, update the state to match your storage system configuration.

1. Boot to maintenance mode:

   ```
   boot_ontap maint
   ```

   a. Enter `y` when you see *Continue with boot?*.

      If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

5. If the displayed system state of the controller does not match your storage system configuration, set the `HA` state for the controller:

   ```
   ha-config modify controller ha
   ```

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed:

   ```
   ha-config show
   ```

Step 2: Verify disk list

1. Verify that the adapter lists the paths to all disks:

   ```
   storage show disk -p
   ```

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode:

```
halt
```

What's next?

After you've restored and verified your system configuration, you need to give back the controller.

Give back the controller - ASA A20, ASA A30, and ASA A50

Return control of storage resources to the replacement controller so your ASA A20, ASA A30, or ASA A50 storage system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption, Onboard Key Manager (OKM) encryption, or External Key Manager (EKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`

   If you encounter errors, contact NetApp Support.

9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

c. Verify status of the keys: `security key-manager key query -restored *false*`

   The output should show no results when when properly synchronized.

   If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode *impaired_node_name*`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback *true*`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

External key manager (EKM)

Reset encryption and return the controller to normal operation.

Steps

1. If the root volume is encrypted with External Key Manager and the console cable is connected to the replacement node, enter `boot_ontap menu` and select option `11`.
2. If these questions appear, answer `y` or `n` as appropriate:

   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}

   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}

   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}

   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}

   Do you know the KMIP server address? {y/n}

   Do you know the KMIP port? {y/n}

   Contact NetApp Support if you have issues.

3. Supply the information for:
   - The client certificate (client.crt) file contents
   - The client key (client.key) file contents
   - The KMIP server CA(s) (CA.pem) file contents
   - The IP address for the KMIP server
   - The port for the KMIP server
4. Once the system processes, you see the Boot Menu. Select '1' for normal boot.
5. Check the takeover status: `storage failover show`
6. Ensure any core dumps on the repaired node are saved by going to advanced mode `set -privilege advanced` and then `run local partner nosavecore`.
7. Return the impaired controller to normal operation by giving back its storage: `storage failover`

```
giveback -ofnode impaired_node_name
```

8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
9. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

After you've transferred the ownership of storage resources to the replacement controller, you need to complete the controller replacement procedure.

Complete controller replacement - ASA A20, ASA A30, and ASA A50

To complete the controller replacement for your ASA A20, ASA A30, or ASA A50 storage system, first restore the NetApp Storage Encryption configuration (if necessary) and install the required licenses on the new controller. Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, register the new controller's serial number and then return the failed part to NetApp.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

  Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The licenses keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs, register the serial number, and check cluster health

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - ASA A20, ASA A30, and ASA A50

Replace a DIMM in your ASA A20, ASA A30, or ASA A50 storage system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

Before you begin

- Ensure all other components in the storage system are working correctly; if not, contact NetApp Support before continuing."

- You must replace the failed FRU component with a replacement FRU component you received from your provider.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

   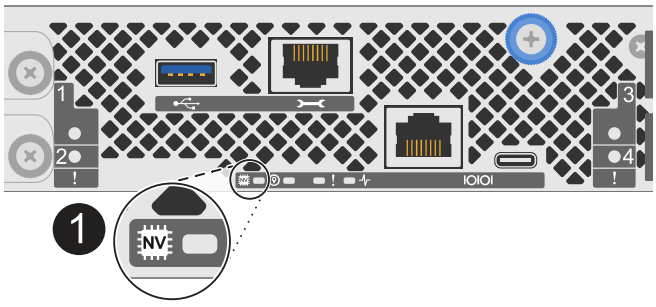

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

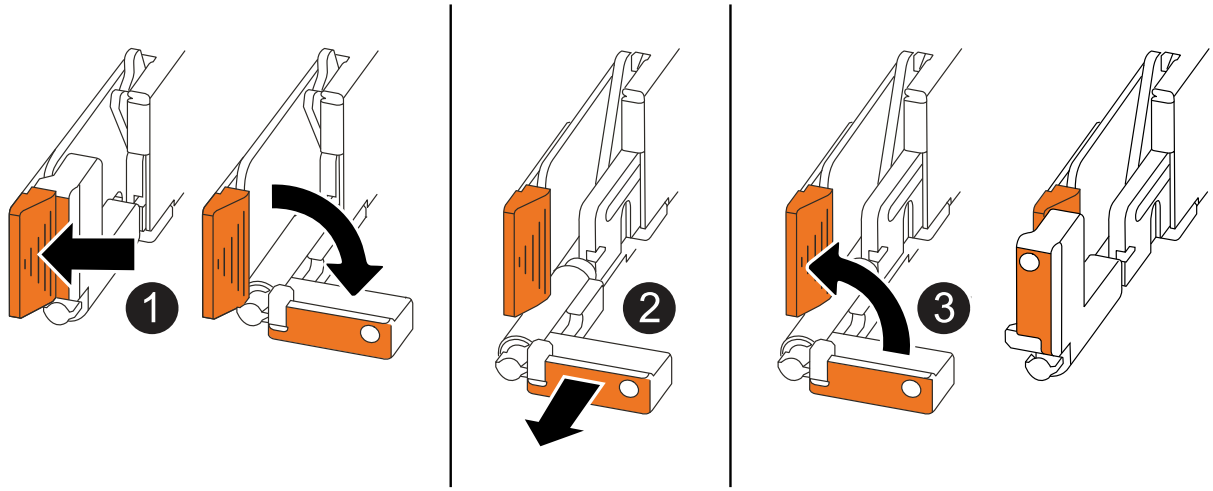

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace a DIMM

To replace a DIMM, locate the faulty DIMM inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller and identify the faulty DIMM.

> Consult either the Netapp Hardware Universe or the FRU map on the cover of the controller for exact DIMM locations.

3. Remove the faulty DIMM:

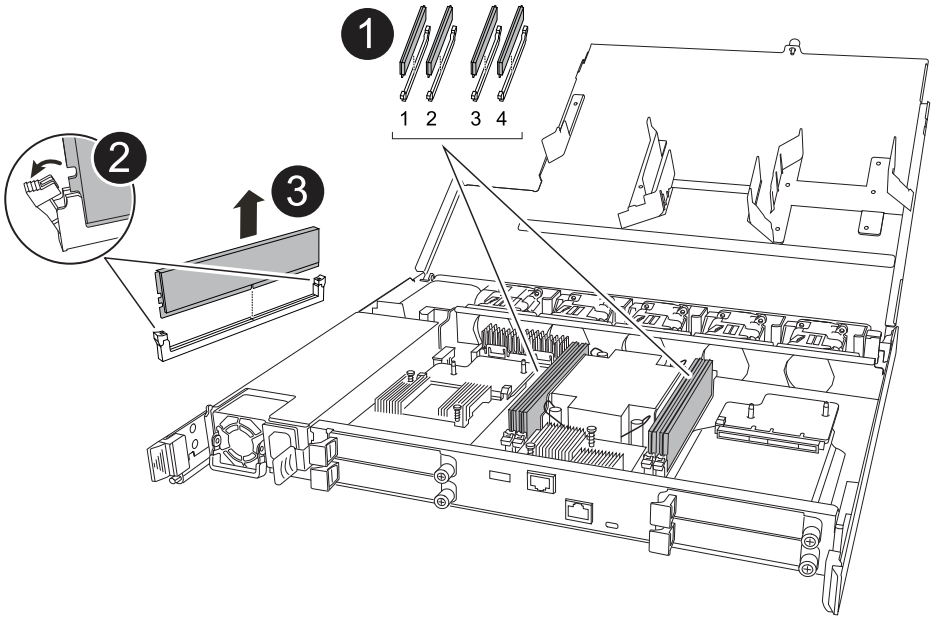

1	DIMM slot numbering and positions. Depending on your storage system model you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM using the same orientation. • Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

4. Install the replacement DIMM:
    a. Remove the replacement DIMM from its antistatic shipping bag.
    b. Make sure that the DIMM ejector tabs on the connector are in the open position.
    c. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

        The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

        When inserted correctly, the DIMM goes in easily but fits tightly in the slot. Reinsert the DIMM if you feel it is not inserted correctly.

    d. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
    e. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

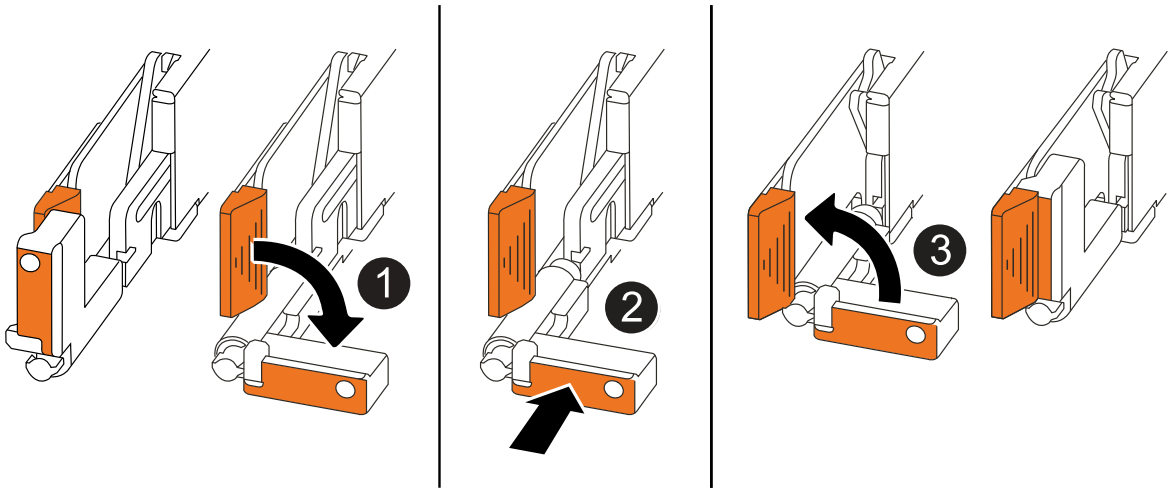

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

> Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

> Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

> Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

> The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a Drive - ASA A20, ASA A30, and ASA A50

Replace a drive in your ASA A20, ASA A30, or ASA A50 storage system when a drive fails or requires an upgrade. The replacement process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed drive nondisruptively while I/O is in progress.

Before you begin

- The drive that you are installing must be supported by your storage system.

  NetApp Hardware Universe

- If self-encrypting drive (SED) authentication is enabled, you must use the SED replacement instructions in the ONTAP documentation.

  Instructions in the ONTAP documentation describe additional steps you must perform before and after replacing an SED.

  NetApp encryption overview with the CLI

- All other components in the storage system must be functioning properly, if not, you must contact NetApp Support before continuing with this procedure.
- Verify that the drive you are removing is failed.

  You can verify that the drive is failed by running the `storage disk show -broken` command. The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  Depending on the drive type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

About this task

- When replacing a failed drive, you must wait 70 seconds between the removal of the drive and the insertion of the replacement drive to allow the storage system to recognize that a drive was removed.
- The best practice is to have the current version of the Disk Qualification Package (DQP) installed before hot-swapping a drive.

  Having the current version of the DQP installed allows your system to recognize and use newly qualified drives. This avoids system event messages about having noncurrent drive information and prevention of drive partitioning because drives are not recognized. The DQP also notifies you of noncurrent drive firmware.

  NetApp Downloads: Disk Qualification Package

- The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on

your system before replacing FRU components.

NetApp Downloads: Disk Shelf Firmware

NetApp Downloads: Disk Drive Firmware

> Do not revert firmware to a version that does not support your shelf and its components.

- Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

  > Drive firmware checks occur every two minutes.

- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Steps

1. Properly ground yourself.
2. Remove the bezel from the front of the storage system.
3. Physically identify the failed drive.
   - When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the shelf operator display panel and the failed drive illuminate.
   - The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:

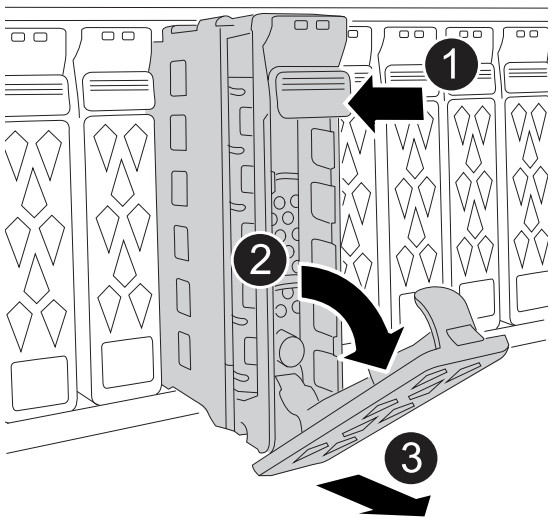

1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Gently push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

7. Verify that the drive's activity (green) LED is illuminated.

   When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat Steps 3 through 7.
9. Reinstall the bezel on the front of the storage system.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support NetApp Support if you need the RMA number or additional help with the replacement procedure.

Replace a fan module - ASA A20, ASA A30, and ASA A50

Replace a fan module in your ASA A20, ASA A30, or ASA A50 storage system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

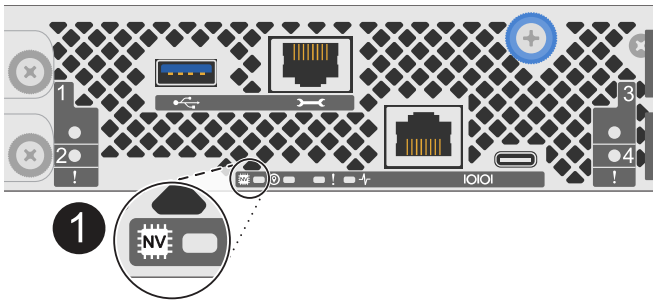

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

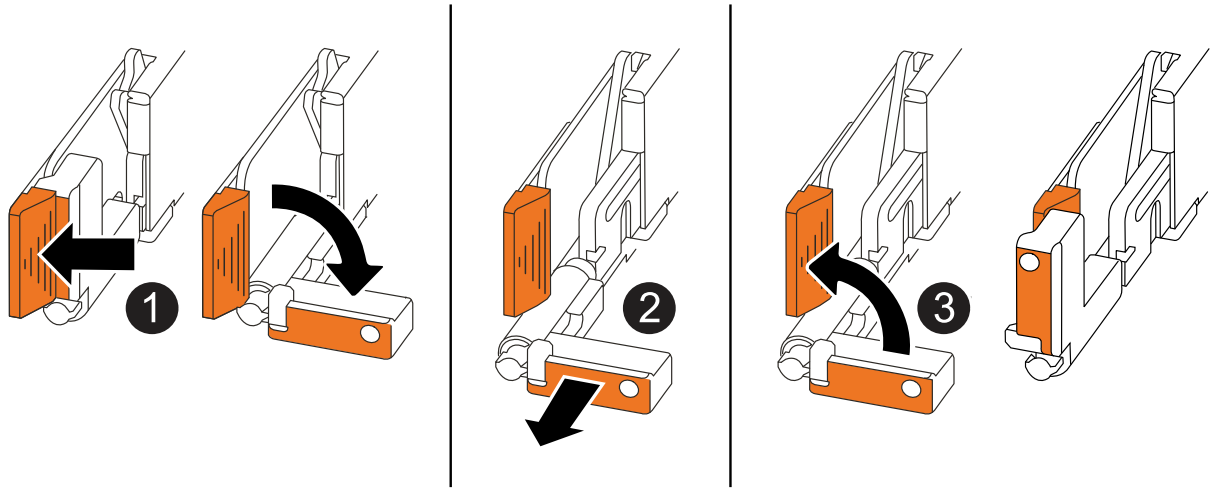

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace fan

To replace a fan, remove the failed fan and replace it with a new fan.

Steps

1. Identify the fan that you must replace by checking the console error messages.
2. Remove the failed fan:

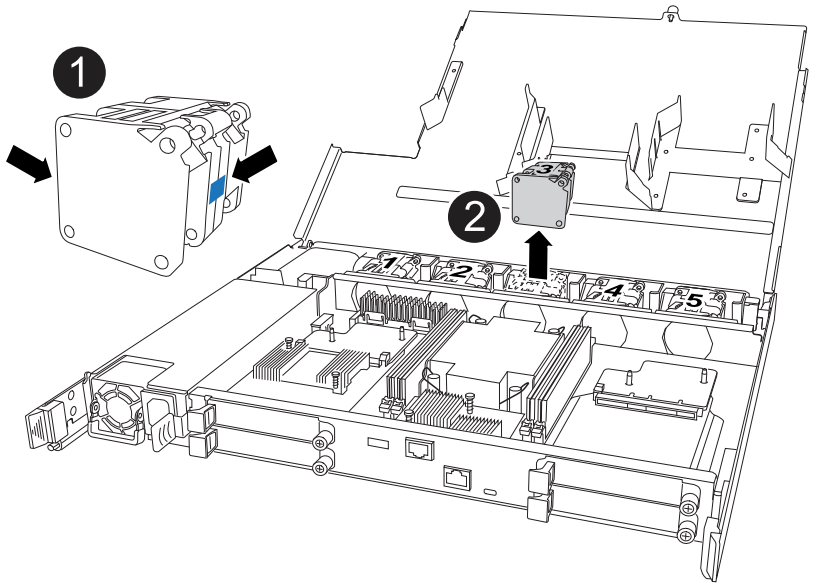

1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

3. Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

Step 4: Reinstall the controller module

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

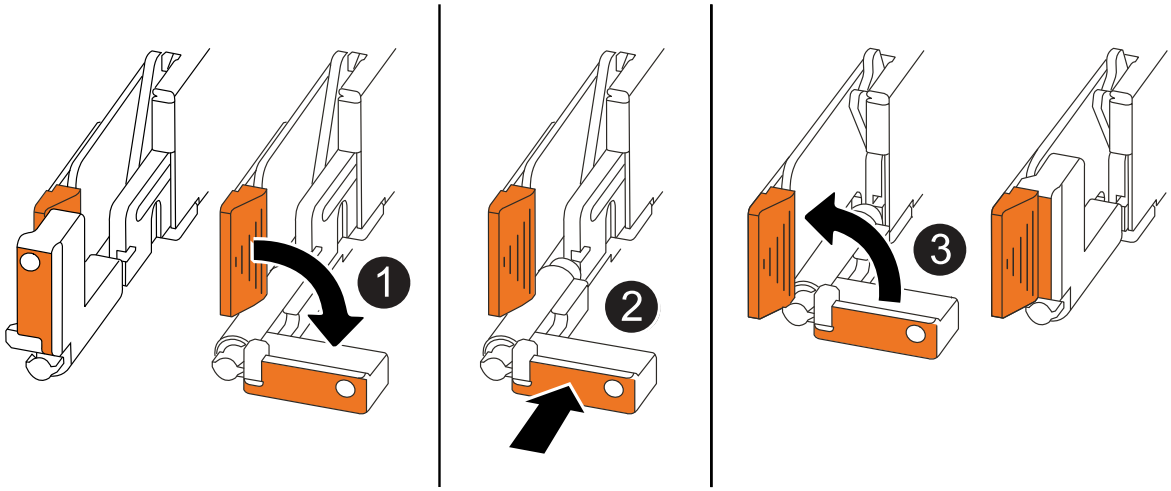

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```

8. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Overview of I/O module maintenance - ASA A20, ASA A30, and ASA A50

The ASA A20, ASA A30, and ASA A50 storage systems offer flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding, hot-swapping or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your storage system with the same type of I/O module, or with a different type of I/O module. You can hot-swap a cluster and HA I/O module when your storage system meets specific requirements. You can also add an I/O module to a storage system with available slots.

- Add an I/O module

  Adding additional I/O modules can improve redundancy, helping to ensure that the storage system remains operational even if one I/O module fails.

- Hot-swap a cluster and HA I/O module

Hot-swapping a failed cluster and HA I/O module can restore the storage system to its optimal operating state. Hot-swapping is done without having to manually take over the impaired controller.

To use this procedure, your storage system must be running ONTAP 9.17.1 or later and meet specific system requirements.

- Replace an I/O module

  Replacing a failing I/O module can restore the storage system to its optimal operating state.

Add an I/O module - ASA A20, ASA A30, and ASA A50

Add an I/O module to your ASA A20, ASA A30, or ASA A50 storage system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your ASA A20, ASA A30, or ASA A50 storage system if there are available slots. If all slots are fully populated, you can replace an existing module to add a new one.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the asaected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, remove the I/O blanking module from the target slot.

   Unused I/O slots should have blanking module installed to prevent possible thermal issues and for EMC compliance.

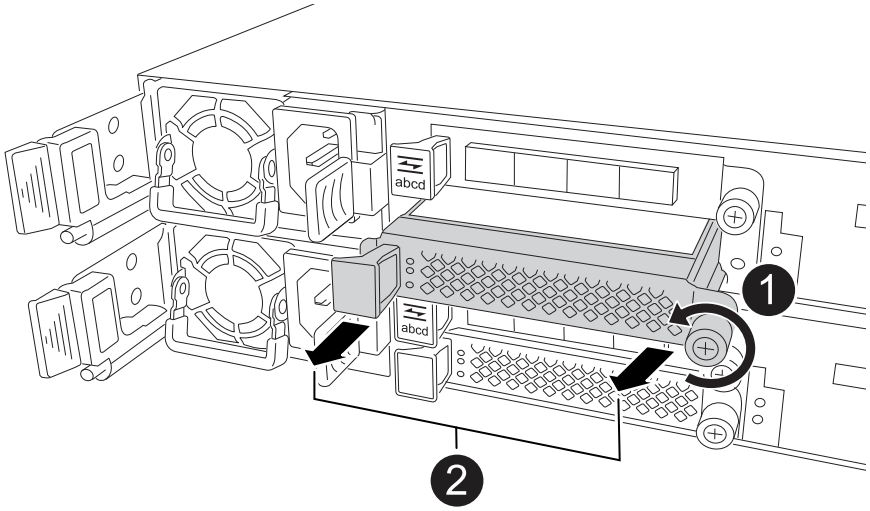

1	On the I/O blanking module, turn the thumbscrew counterclockwise to loosen.
2	Pull the I/O blanking module out of the controller using the tab on the left and the thumbscrew.

3. Install the new I/O module:
   a. Align the I/O module with the edges of the controller slot opening.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O module.

   c. Turn the thumbscrew clockwise to tighten.
4. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.

5. Reboot the impaired controller from the LOADER prompt: `bye`

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

6. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name.
```

7. Repeat these steps to add an I/O module to the other controller.
8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation: +

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, unplug any cabling on the target I/O module.
3. Remove the target I/O module from the controller:

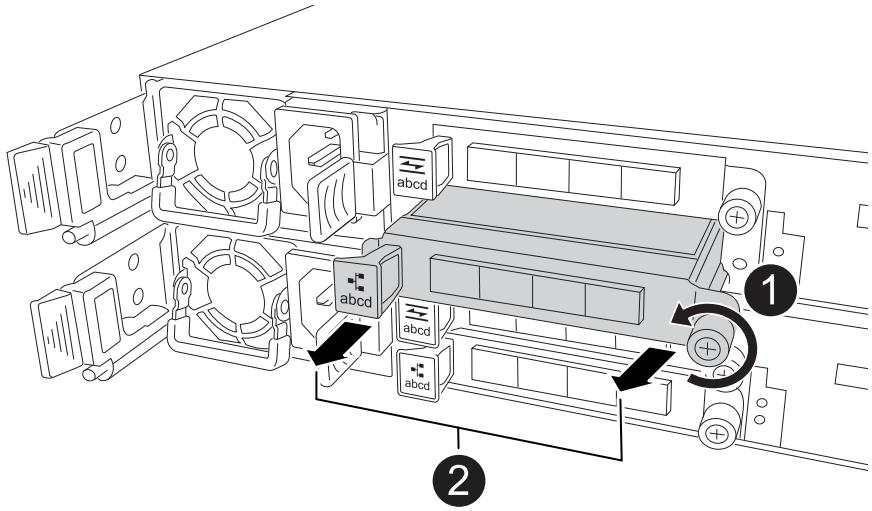

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the new I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
5. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
6. Repeat the I/O module remove and install steps to add any additional I/O modules in the controller.
7. Reboot the impaired controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.
8. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```
9. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```
10. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```
11. If you installed a NIC module, specify the usage mode for each port as *network*:

    ```
    storage port modify -node node_name -port port_name -mode network
    ```
12. Repeat these steps for the other controller.

Hot-swap the I/O module used for cluster and HA traffic - ASA A20, ASA A30, and ASA A50

The cluster and HA I/O module supports interconnects for clustering and high-availability. You can hot-swap the module in your ASA A20, ASA A30, or ASA A50 storage system when the module fails and if your storage system meets specific requirements.

To hot-swap a module, you ensure your storage system meets the procedure requirements, prepare the storage system and I/O module in slot 4, hot-swap the failed module for an equivalent one, bring the replacement module online, restore the storage system to normal operation, and return the failed module to NetApp.

About this task

- Hot-swapping the cluster and HA I/O module means that you do not have to perform a manual takeover; the impaired controller (the controller with the failed cluster and HA I/O module) has automatically taken over the healthy controller.

  When the impaired controller has taken over the healthy controller, the only way to recover without an outage is to hot-swap the module.

- It is critical to apply the commands to the correct controller when you are hot-swapping the cluster and HA I/O module:
  - The *impaired controller* is the controller on which you are hot-swapping the cluster and HA I/O module and it is the controller that has taken over the healthy controller.
  - The *healthy controller* is the HA partner of the impaired controller and it is the controller that was taken over by the impaired controller.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, make sure your storage system meets all requirements.

> If your storage system does not meet all requirements, you must use the replace an I/O module procedure.

- Your storage system must be running ONTAP 9.17.1 or later.
- The I/O module that failed must be a cluster and HA I/O module in slot 4 and you must be replacing it with an equivalent cluster and HA I/O module. You cannot change the I/O module type.
- Your storage system configuration must have only one cluster and HA I/O module located in slot 4, not two cluster and HA I/O modules.
- Your storage system must be a two-node (switchless or switched) cluster configuration.
- The controller with the failed cluster and HA I/O module (the impaired controller) must have already taken over the healthy partner controller. The takeover should have occurred automatically if the I/O module is failed.

  For two-node clusters, the storage system cannot discern which controller has the failed I/O module, so either controller might initiate the takeover. The cluster and HA I/O module hot-swap procedure is only supported when the controller with the failed I/O module (the impaired controller) has taken over the healthy controller.

You can verify that the impaired controller successfully took over the healthy controller by entering the `storage failover show` command.

If you are not sure which controller has the failed I/O module, contact NetApp Support.

- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Step 2: Prepare the storage system and I/O module slot 4

Prepare the storage system and I/O module slot 4 so that it is safe to remove the failed cluster and HA I/O module:

Steps

1. Properly ground yourself.
2. Unplug cabling from the failed cluster and HA I/O module.

   Make sure to label the cables so that later in this procedure you can reconnect them to the same ports.

3. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<number of hours down>h
   ```

   For example, the following AutoSupport message suppresses automatic case creation for two hours:

   ```
   node2::> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

4. Disable automatic giveback:
   a. Enter the following command from the console of the impaired controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
5. Prepare the failed cluster and HA module in slot 4 for removal by removing it from service and powering it off:
   a. Enter the following command:

      ```
      system controller slot module remove -node impaired_node_name -slot slot_number
      ```

   b. Enter `y` when you see the prompt *Do you want to continue?*

      For example, the following command prepares the module in slot 4 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

```
node2::> system controller slot module remove -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered off for removal.

Do you want to continue? {y|n}: y

The module has been successfully removed from service and powered
off. It can now be safely removed.
```

6. Verify the failed cluster and HA module in slot 4 is powered off:

   `system controller slot module show`

   The output should show *`powered-off`* in the status column for the failed module in slot 4.

Step 3: Replace the failed cluster and HA I/O module

Replace the failed cluster and HA I/O module in slot 4 with an equivalent I/O module:

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the failed cluster and HA I/O module from the impaired controller:

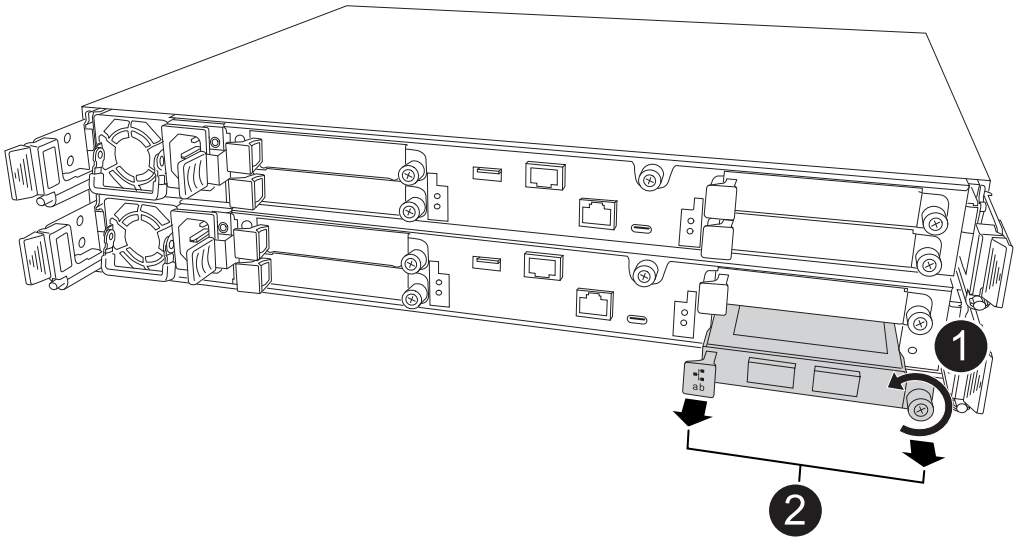

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew on the right.

3. Install the replacement cluster and HA I/O module into slot 4:

a. Align the I/O module with the edges of the slot.

b. Gently push the I/O module all the way into the slot, making sure to properly seat the I/O module into the connector.

   You can use the tab on the left and the thumbscrew on the right to push in the I/O module.

c. Turn the thumbscrew clockwise to tighten.

4. Cable the cluster and HA I/O module.

Step 4: Bring the replacement cluster and HA I/O module online

Bring the replacement cluster and HA I/O module in slot 4 online, verify the module ports initialized successfully, verify slot 4 is powered on, and then verify the module is online and recognized.

Steps

1. Bring the replacement cluster and HA I/O module online:

   a. Enter the following command:

      `system controller slot module insert -node` *`impaired_node_name`* `-slot` *`slot_name`*

   b. Enter `y` when you see the prompt, *Do you want to continue?*

      The output should confirm the cluster and HA I/O module was successfully brought online (powered on, initialized, and placed into service).

      For example, the following command brings slot 4 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered on and initialized.

Do you want to continue? {y|n}: `y`

The module has been successfully powered on, initialized and placed
into service.
```

2. Verify that each port on the cluster and HA I/O module successfully initialized:

   `event log show -event *hotplug.init*`

   > It might take several minutes to allow for any required firmware updates and port initialization.

   The output should show a hotplug.init.success EMS event logged for each port on the cluster and HA I/O module with *`hotplug.init.success:`* in the *`Event`* column.

For example, the following output shows initialization succeeded for cluster and HA I/O module ports e4b and e4a:

```
node2::> event log show -event *hotplug.init*

Time                Node             Severity      Event

------------------- ---------------- -------------
---------------------------

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4b" in slot 4 succeeded

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4a" in slot 4 succeeded

2 entries were displayed.
```

3. Verify I/O module slot 4 is powered on and ready for operation:

   `system controller slot module show`

   The output should show slot 4 status as *`powered-on`* and therefore ready for operation of the replacement cluster and HA I/O module.

4. Verify that the replacement cluster and HA I/O module is online and recognized.

   Enter the command from the console of the impaired controller:

   `system controller config show -node local -slot4`

   If the replacement cluster and HA I/O module was successfully brought online and is recognize, the output shows I/O module information, including port information, for slot 4.

   For example, you should see output similar to the following:

```
node2::> system controller config show -node local -slot 4

Node: node2
Sub- Device/
Slot slot Information
---- ---- ------------------------------
   4    - Dual 40G/100G Ethernet Controller CX6-DX
                  e4a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2807GJFM-B
                  e4b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2809G26F-A
                  Device Type:        CX6-DX PSID(NAP0000000027)
                  Firmware Version:   22.44.1700
                  Part Number:        111-05341
                  Hardware Revision:  20
                  Serial Number:      032403001370
```

Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the healthy controller, restoring automatic giveback, and reenabling AutoSupport automatic case creation.

Steps

1. Return the healthy controller (the controller that was taken over) to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode healthy_node_name
   ```

2. Restore automatic giveback from the console of the impaired controller (the controller that took over the healthy controller):

   ```
   storage failover modify -node local -auto-giveback true
   ```

3. If AutoSupport is enabled, restore automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=end
   ```

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an I/O module - ASA A20, ASA A30, and ASA A50

Replace an I/O module in your ASA A20, ASA A30, or ASA A50 storage system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

Use this procedure to replace a failed I/O module.

Before you begin

All other components in the storage system must be functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace a failed I/O module

To replace a failed I/O module, locate it in the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug cabling from the failed I/O module.

   Make sure to label the cables so that you know where they came from.

3. Remove the failed I/O module from the controller:

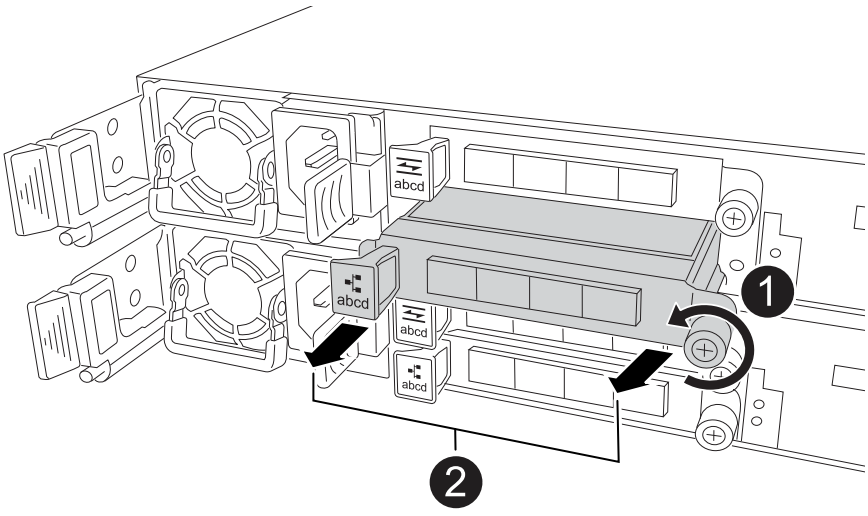

1	Turn the I/O module thumbscrew counterclockwise to loosen.

2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the replacement I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
5. Cable the I/O module.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt: `bye`

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the node to normal operation: `storage failover giveback -ofnode` *`impaired_node_name`*
3. Restore automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback` *`true`*

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NV battery - ASA A20, ASA A30, and ASA A50

Replace the NV battery in your ASA A20, ASA A30, or ASA A50 storage system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

To replace the NV battery, you must remove the controller, remove the faulty battery, install the replacement battery, and then reinstall the controller.

Before you begin

All other components in the storage system must be functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected

storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

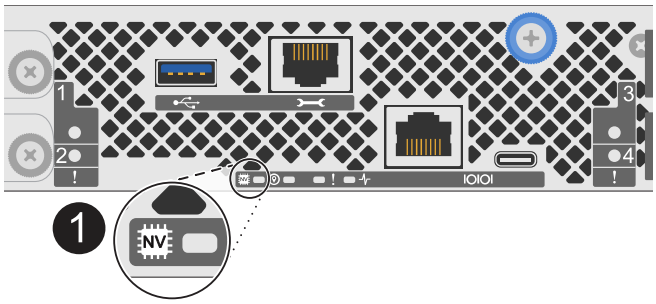

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

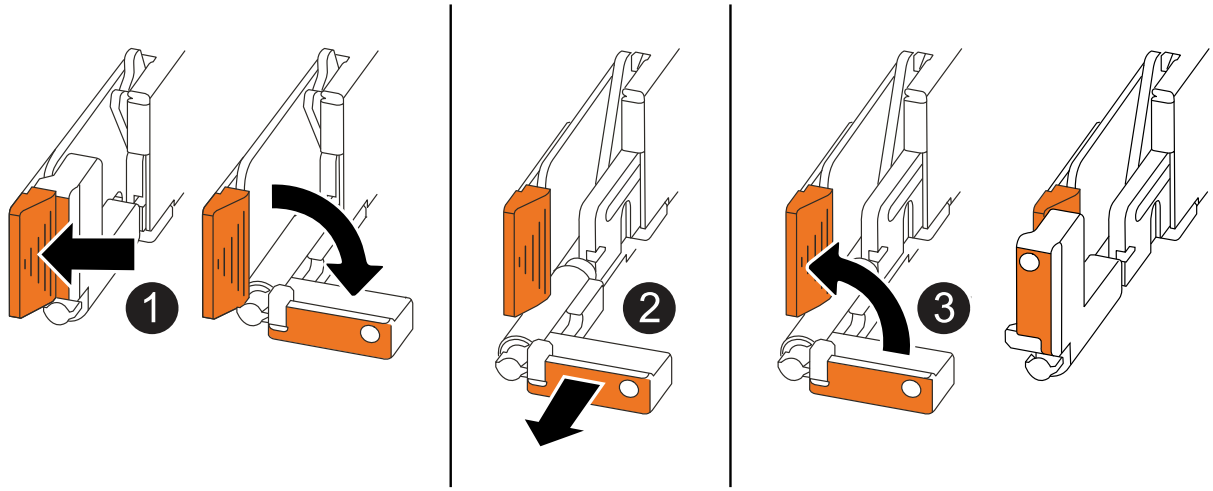

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the NV battery

Remove the failed NV battery from the controller and install the replacement NV battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the NV battery.
3. Remove the NV battery:

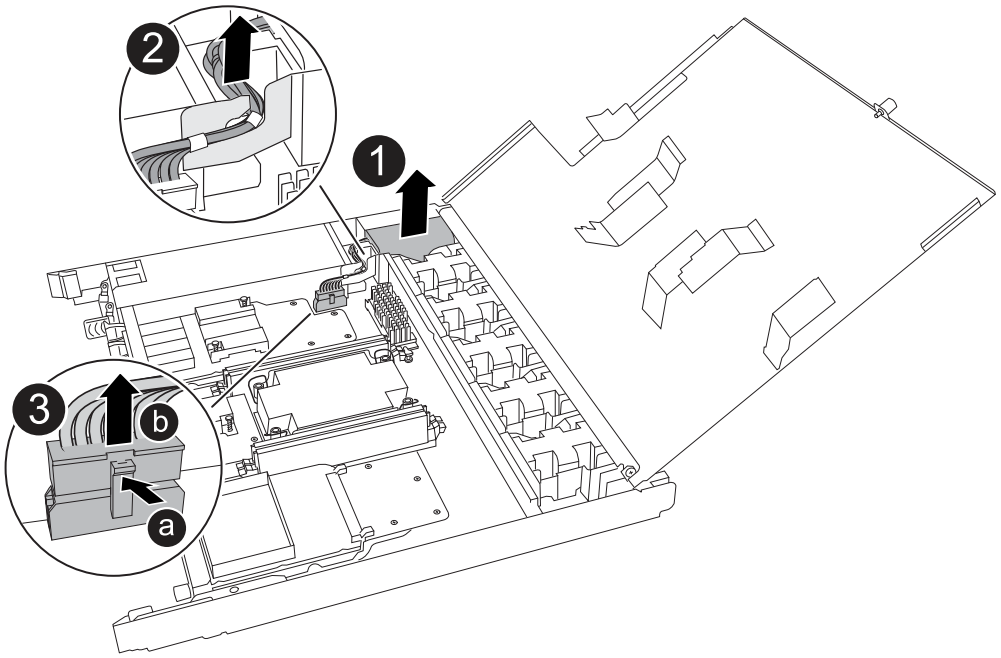

1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

4. Install the replacement NV battery:
   a. Remove the replacement battery from its package.
   b. Plug the wiring connector into its socket.
   c. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   d. Place the NV battery into its compartment.

      The NV battery should sit flush in its compartment.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

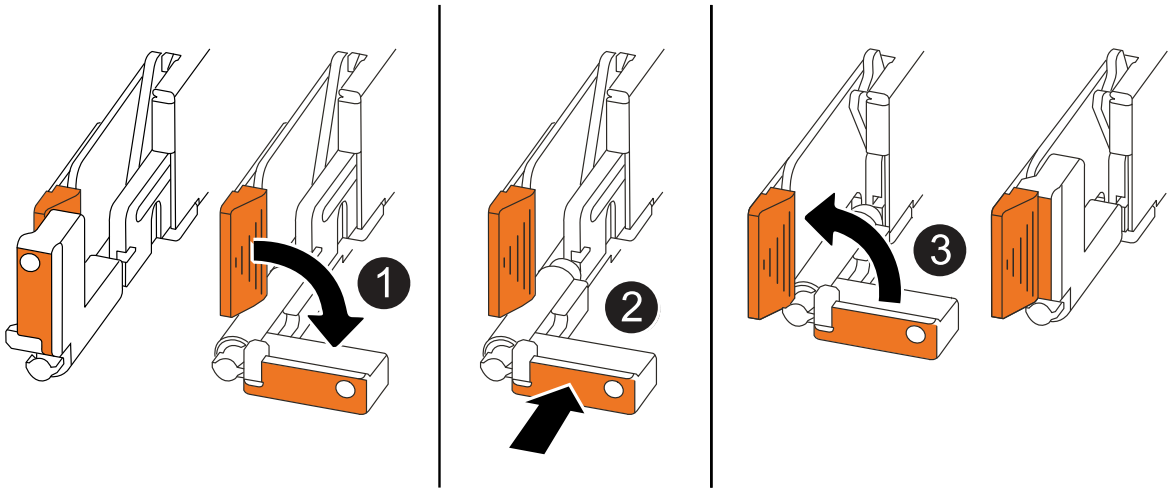

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.

The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - ASA A20, ASA A30, and ASA A50

Replace an AC or DC power supply unit (PSU) in your ASA A20, ASA A30, or ASA A50 storage system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cord, replacing the faulty PSU, and then reconnecting it to the power source.

About this task

- This procedure is written for replacing one PSU at a time.

The PSUs are redundant and hot-swappable.

- IMPORTANT: Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.
- Use the appropriate procedure for your type of PSU: AC or DC.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the power cord from the PSU by opening the power cord retainer, and then unplug the power cord from the PSU.

   PSUs do not have a power switch.

4. Remove the PSU:

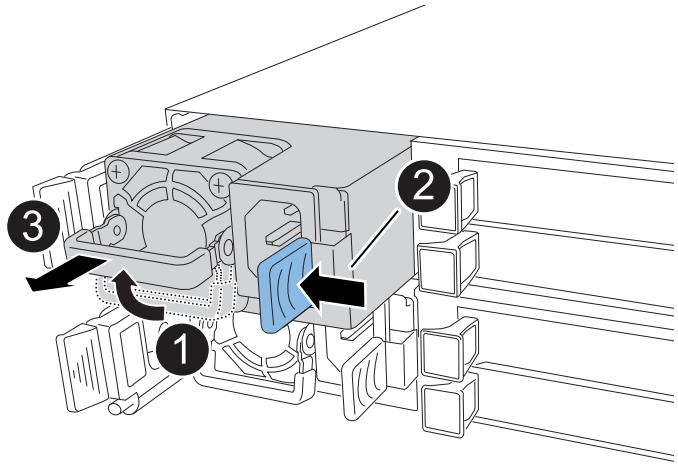

1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

5. Install the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

c. Rotate the handle down, so it is out of the way of normal operations.

6. Reconnect the power cord to the PSU and secure the power cord with the power cord retainer.

   Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the PSU:

   PSUs do not have a power switch.

   a. Unscrew the two thumb screws on the D-SUB DC power cord connector.

      The illustration and table in step 4 shows the two thumb screws (item #1) and the D-SUB DC power cord connector (item #2).

   b. Unplug the cord from the PSU and set it aside.

4. Remove the PSU:

   a. Rotate the handle up, to its horizontal position, and then grasp it.

   b. With your thumb, press the terracotta tab to release the locking mechanism.

   c. Pull the PSU out of the controller while using your other hand to support its weight.

   The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.

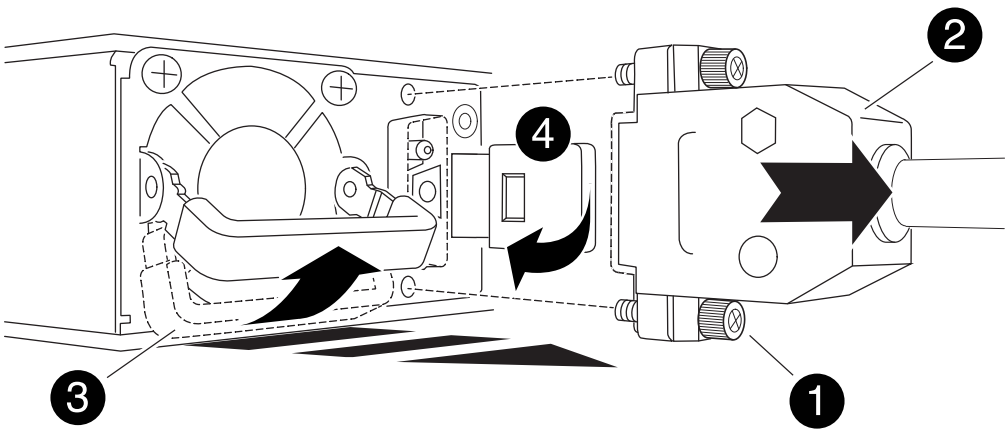

1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

5. Insert the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently slide the PSU into the controller until the locking tab clicks into place.

      A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.
6. Reconnect the D-SUB DC power cord:

   Once power is restored to the PSU, the status LED should be green.

   a. Plug the D-SUB DC power cord connector into the PSU.
   b. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA A20, ASA A30, and ASA A50

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your ASA A20, ASA A30, or ASA A50 storage system to ensure that services and applications relying on accurate time synchronization remain operational.

You replace the real-time clock (RTC) battery in the controller so that your storage system's services and applications that depend on accurate time synchronization continue to function.

Before you begin

All other components in the storage system must be functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

About this task

- You can use this procedure with all versions of ONTAP supported by your storage system.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

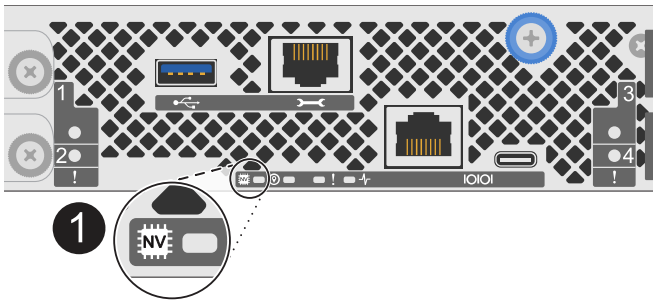

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

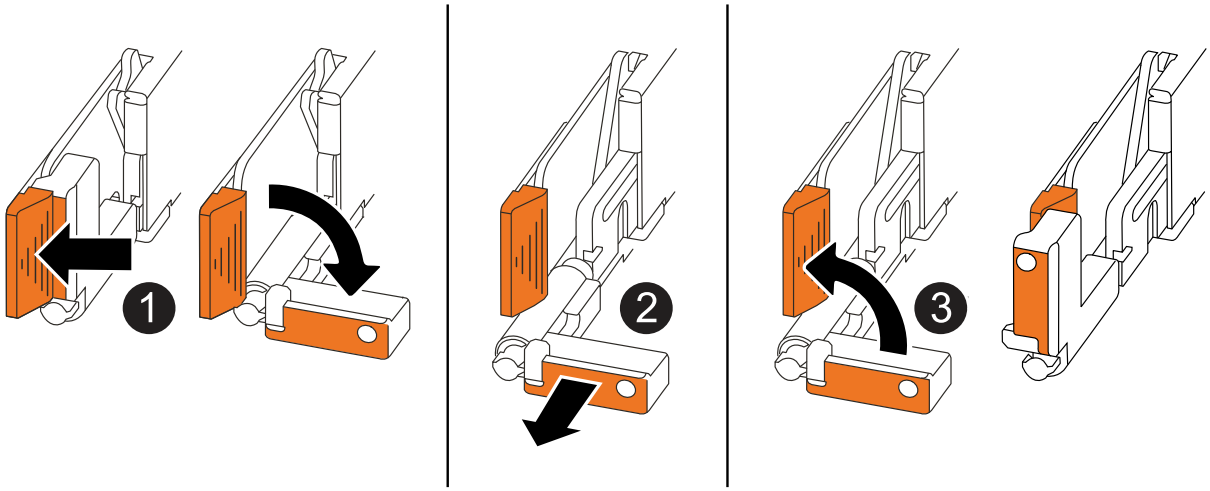

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the RTC battery

Remove the failed RTC battery and install the replacement RTC battery.

Steps

1. Locate the RTC battery.
2. Remove the RTC battery:

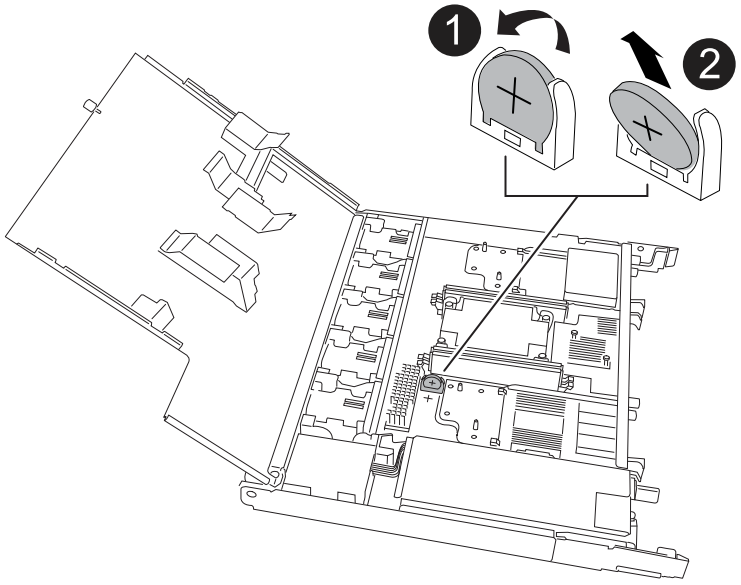

1	Gently rotate the RTC battery at an angle away from its holder.
2	Lift the RTC battery out of its holder.

3. Install the replacement RTC battery:
   a. Remove the replacement battery from the antistatic shipping bag.
   b. Position the battery so that the plus sign on the battery faces out to correspond with the plus sign on the motherboard.
   c. Insert the battery into the holder at an angle, and then push it into an upright position so it is fully seated in the holder.
   d. Visually inspect the battery to make sure that it is completely seated in its holder and that the polarity is correct.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

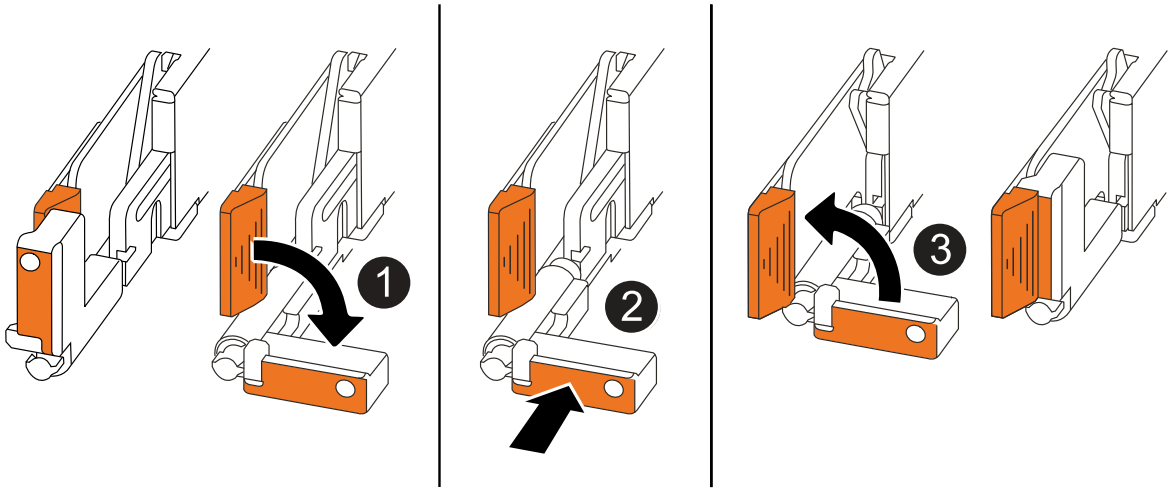

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode` *`impaired_node_name`*

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Reset the time and date on the controller

After replacing the RTC battery, inserting the controller and powering on first BIOS reset, you will see the following error messages:
`RTC date/time error. Reset date/time to default`
`RTC power failure error`
These messages are expected and you can continue with this procedure.

1. On the healthy controller, check the date and time with the `cluster date show` command.

   If your storage system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*.

2. On the impaired controller, at the LOADER prompt, check the time and date: `cluster date show`
   a. If necessary, modify the date: `set date` *`mm/dd/yyyy`*
   b. If necessary, set the time, in GMT: `set time hh:mm:ss`
   c. Confirm the date and time.
3. At the LOADER prompt, enter `bye` to reinitialize the I/O modules, other components, and let the controller reboot.
4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

ASA C30 systems

Overview of hardware maintenance - ASA C30

Maintain the hardware of your ASA C30 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The procedures in this section assume that the ASA C30 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the ASA C30 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller and drives.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
NV battery	The non-volatile memory (NV) battery is responsible for providing power to the NVMEM components while data in-flight is being destaged to flash memory after a power loss.

Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.

Boot media

Boot media replacement workflow - ASA C30

Get started with replacing the boot media in your ASA C30 storage system by reviewing the replacement requirements, shutting down the impaired controller, replacing the boot media, restoring the image on the boot media, and verifying the system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the impaired controller**

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

3. **Replace the boot media**

Remove the failed boot media from the impaired controller and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the healthy controller.

5. **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements to replace the boot media - ASA C30

Before replacing the boot media in your ASA C30 storage system, ensure you meet the necessary requirements and considerations for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0M (wrench) port on the impaired controller is working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller to replace the boot media - ASA C30

Shut down the impaired controller in your ASA C30 storage system to prevent data loss and ensure system stability when replacing the boot media.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media - ASA C30

The boot media in your ASA C30 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then transferring the ONTAP image to the replacement boot media.

About this task

If needed, you can turn on the platform chassis location (blue) LEDs to aid in physically locating the affected platform. Log into the BMC using SSH and enter the `system location-led` *`on`* command.

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led` *`off`* command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

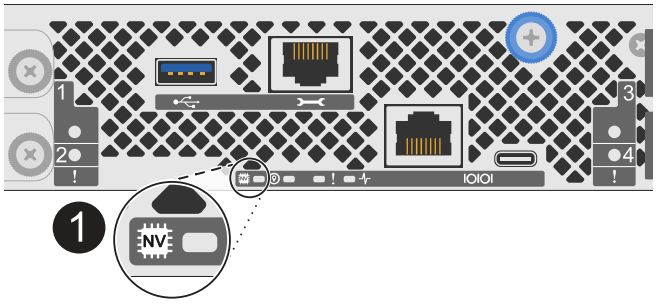

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

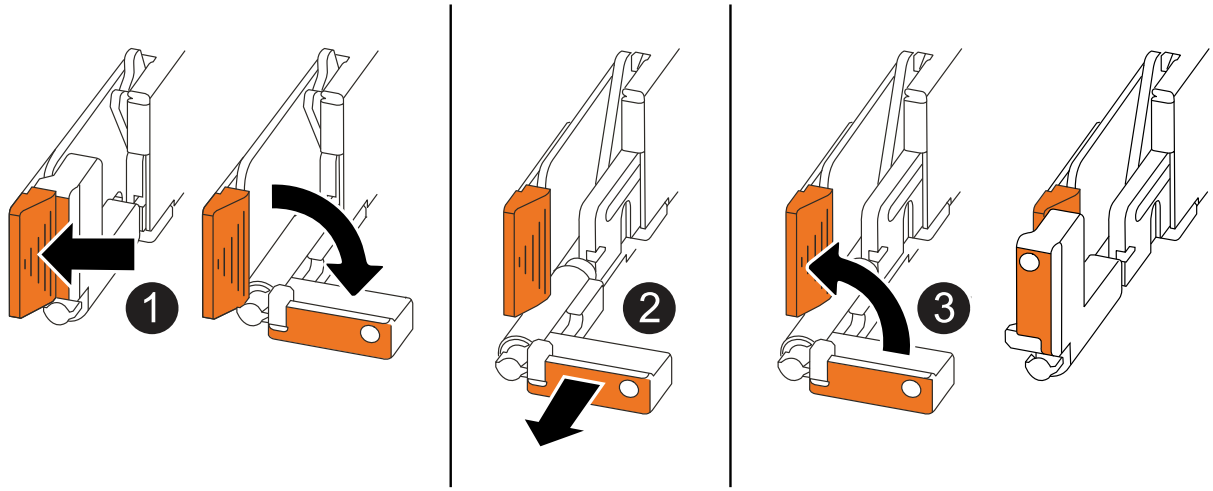

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.
7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:

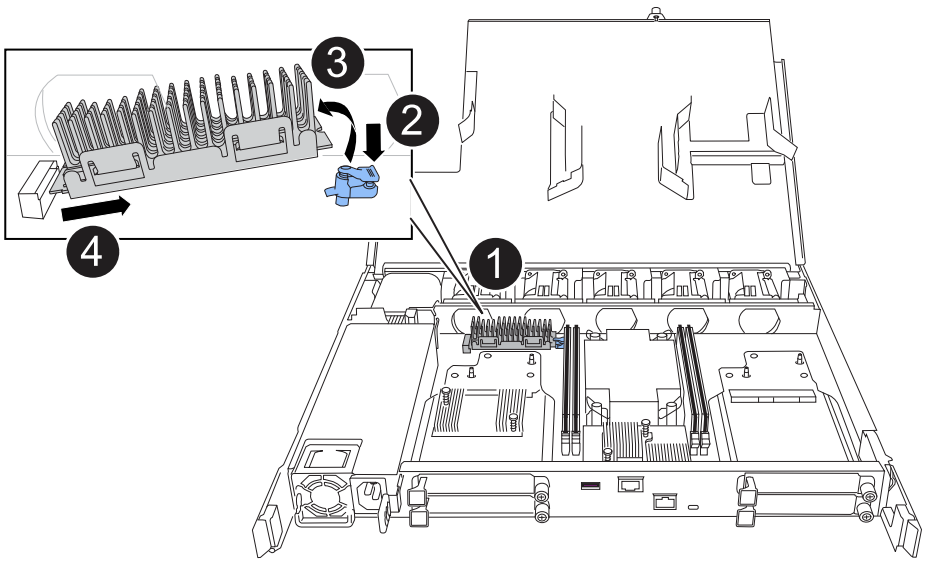

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:
   a. Remove the boot media from its package.
   b. Slide the socket end of the boot media into its socket.
   c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

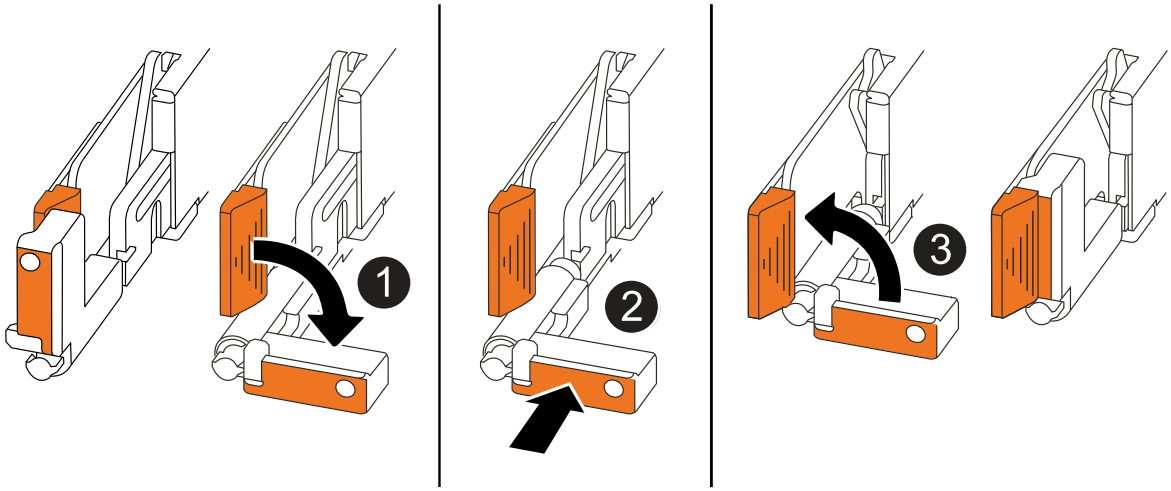

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   > Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.

   > Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

      > The controller boots to the LOADER prompt when fully seated in the chassis. It gets its power from the partner controller.

b. Rotate the controller handles up and lock in place with the tabs.

5. Reconnect the power cord to the PSU on the impaired controller.

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next

After physically replacing the impaired boot media, you restore the ONTAP image from the partner node.

Restore the ONTAP image on the boot media - ASA C30

After installing the new boot media device in your ASA C30 storage system, you can start the automated boot media recovery process to restore the configuration from the healthy node.

During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed part to NetApp - ASA C30

When a component in your ASA C30 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - ASA C30

Replacing the chassis in your ASA C30 storage system consists of reviewing the replacement requirements, shutting down the controllers, replacing the chassis, and verifying system operations.

Review the chassis replace requirements

Review the requirements to replace the chassis.

Shut down the controllers

Shut down the controllers so you can perform maintenance on the chassis.

Replace the chassis

Replace the chassis by moving the drives and any drive blanks, controllers (with the power supplies), and bezel from the impaired chassis to the new chassis, and swapping out the impaired chassis with the new chassis of the same model as the impaired chassis.

Complete chassis replacement

Verify the HA state of the chassis and return the failed part to NetApp.

Requirements to replace the chassis - ASA C30

Before replacing the chassis of your ASA C30 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement chassis, and the necessary tools.

Review the following requirements and considerations.

Requirements

- The replacement chassis must be the same model as the impaired chassis. This procedure is for a like-for-like replacement, not for an upgrade.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Considerations

- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your storage system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, drives, any drive blanks, and controllers to the new chassis.

What's next?

After you've reviewed the requirements to replace the chassis, you need to shut down the controllers

Shut down the controllers - ASA C30

Shut down the controllers in your ASA C30 storage system to prevent data loss and ensure system stability when replacing the chassis.

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

```
system node autosupport invoke -node * -type all -message "MAINT=2h Replace
```

chassis"

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

What's next?

After you've shut down the controllers, you need to replace the chassis.

Replace the chassis - ASA C30

Replace the chassis of your ASA C30 storage system when a hardware failure requires it. The replacement process involves removing the controllers, removing the drives, installing the replacement chassis, and reinstalling the chassis components.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

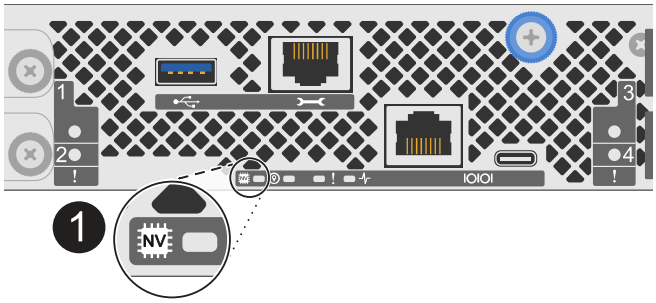

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

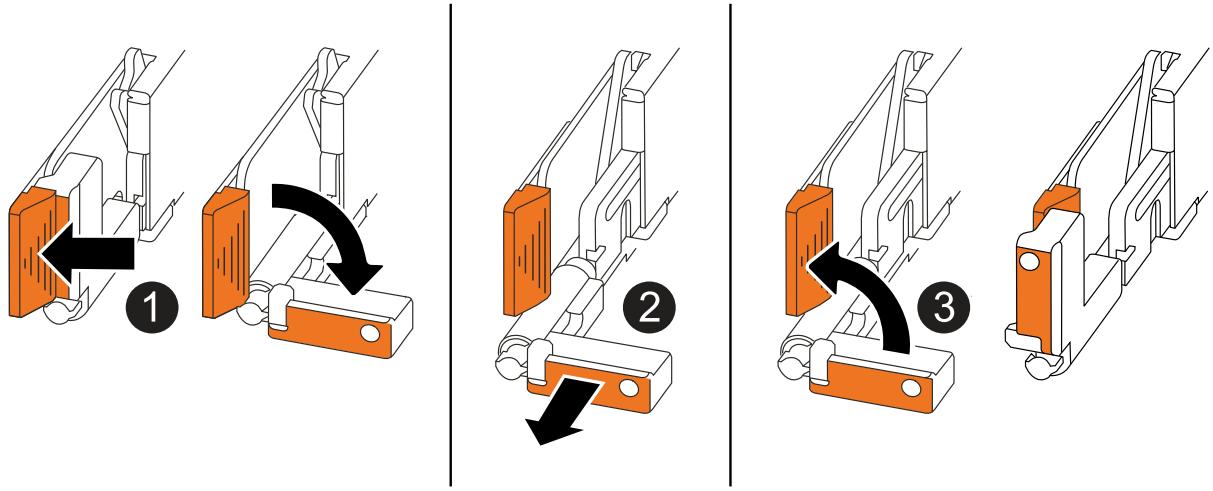

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Repeat these steps for the other controller in the chassis.

Step 2: Remove the drives from the impaired chassis

You need to remove all of the drives and any drive blanks from the impaired chassis so that later in the procedure you can install them in the replacement chassis.

1. Gently remove the bezel from the front of the storage system.
2. Remove the drives and any drive blanks:

   Keep track of what drive bay each drive and drive blank was removed from because they must be installed in the same drive bays in the replacement chassis.

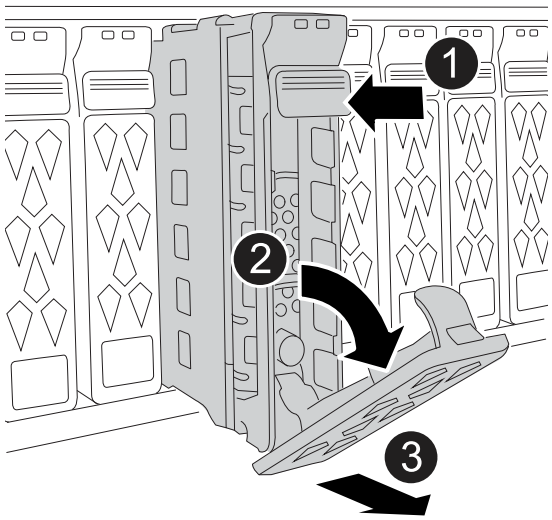

1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

3. Set the drives aside on a static-free cart or table.

Step 2: Replace the chassis from within the equipment rack or system cabinet

You remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis, install the drives, any drive blanks, and then install the bezel.

1. Remove the screws from the impaired chassis mount points.

   Set the screws aside to use later in this procedure.

   If the storage system shipped in a NetApp system cabinet, you must remove additional screws at the rear of the chassis before the chassis can be removed.

2. Using two people or a power lift, remove the impaired chassis from the equipment rack or system cabinet by sliding it off the rails, and then set it aside.
3. Using two people, install the replacement chassis into the equipment rack or system cabinet by sliding it onto the rails.
4. Secure the front of the replacement chassis to the equipment rack or system cabinet using the screws you removed from the impaired chassis.

Step 4: Install the controllers and drives

Install the controllers and drives into the replacement chassis and reboot the controllers.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when installing a controller, and can be used as a reference for the rest of the controller installation steps.

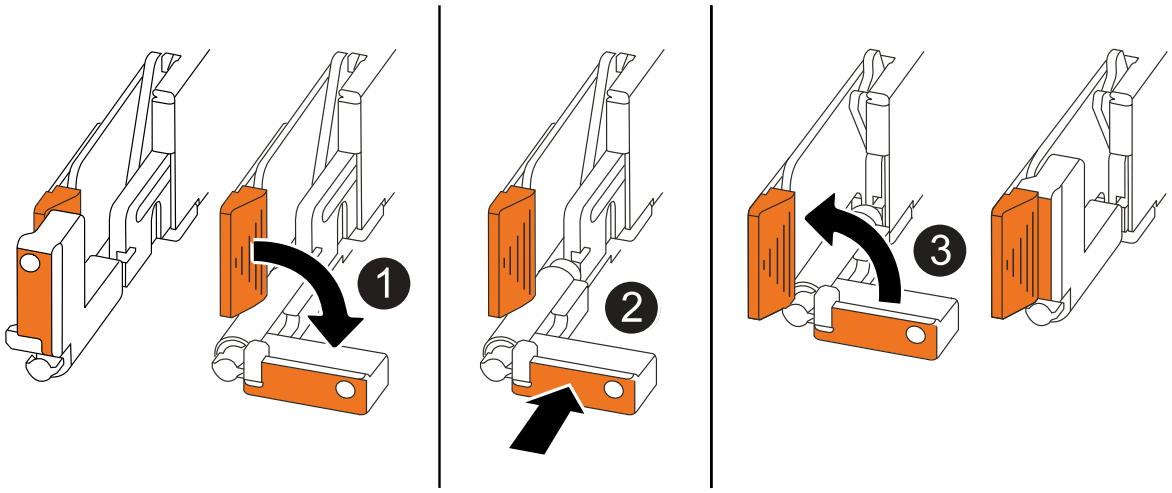

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis and push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

1. Insert one of the controllers into the chassis:
   a. Align the back of the controller with the opening in the chassis.
   b. Firmly push on the handles until the controller meets the midplane and is fully seated in the chassis.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   c. Rotate the controller handles up and lock in place with the tabs.
2. Recable the controller, as needed, except for the power cords.
3. Repeat these steps to install the second controller into the chassis.
4. Install the drives and any drive blanks you removed from the impaired chassis into the replacement chassis:

   The drives and drive blanks must be installed in the same drive bays in the replacement chassis.

   a. With the cam handle in the open position, use both hands to insert the drive.

b. Gently push until the drive stops.

c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

d. Repeat the process for the remaining drives.

5. Install the bezel.
6. Reconnect the power cords to the power supplies (PSU) in the controllers.

   Once power is restored to a PSU, the status LED should be green.

   The controllers begin to boot as soon as the power is restored.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. If controllers boot to the LOADER prompt, reboot the controllers:

```
boot_ontap
```

8. Turn AutoSupport back on:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

After you've replaced the impaired chassis and reinstalled the components into it, you need to complete the chassis replacement.

Complete chassis replacement - ASA C30

Verify the HA state of the chassis and then return the failed part to NetApp to complete the final step in the ASA C30 chassis replacement procedure.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your storage system configuration.

1. In Maintenance mode, from either controller, display the HA state of the local controller and chassis:

```
ha-config show
```

The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your storage system configuration:
   a. Set the HA state for the chassis:

      `ha-config modify chassis HA-state`

      The value for HA-state should be *ha*.
      The value for HA-state can be one of the following:
      * ***ha***
      * *mcc* (not supported in ASA)

   b. Confirm that the setting has changed:

      `ha-config show`

3. If you have not already done so, recable the rest of your storage system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - ASA C30

Get started with replacing the controller in your ASA C30 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

Review the controller replacement requirements

To replace the controller, you must meet certain requirements.

Shut down the impaired controller

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

Replace the controller

Replacing the controller includes removing the impaired controller, moving FRU components to the replacement controller, installing the replacement controller in the chassis, setting the time and date, and then recabling.

Restore and verify the system configuration

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

5 Give back the controller

Transfer the ownership of storage resources back to the replacement controller.

6 Complete controller replacement

Verify the LIFs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - ASA C30

Before replacing the controller in your ASA C30 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

You must review the requirements and considerations for the controller replacement procedure.

Requirements

- All shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- You must replace a controller with a controller of the same model type. You cannot upgrade your system by just replacing the controller.
- You cannot change any drives or shelves as part of this procedure.
- You must always capture the controller's console output to a text log file.

  The console output provides you with a record of the procedure you can use to troubleshoot issues you might encounter during the replacement process.

Considerations

- It is important that you apply the commands in this procedure to the correct controller:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.

What's next?

After you've reviewed the requirements to replace the impaired controller, you need to shut down the impaired controller.

Shut down the impaired controller - ASA C30

Shut down the impaired controller in your ASA C30 storage system to prevent data loss and ensure system stability when replacing the controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After you've shut down the impaired controller, you need to replace the controller.

Replace the controller - ASA C30

Replace the controller in your ASA C30 storage system when a hardware failure requires

it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

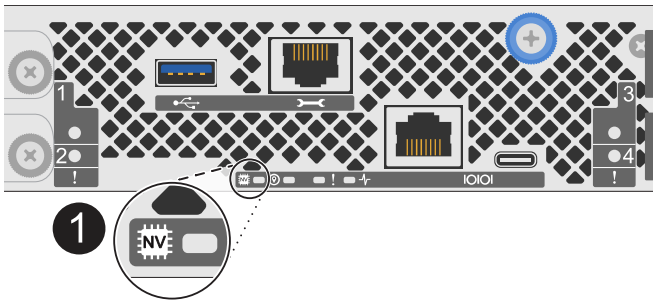

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

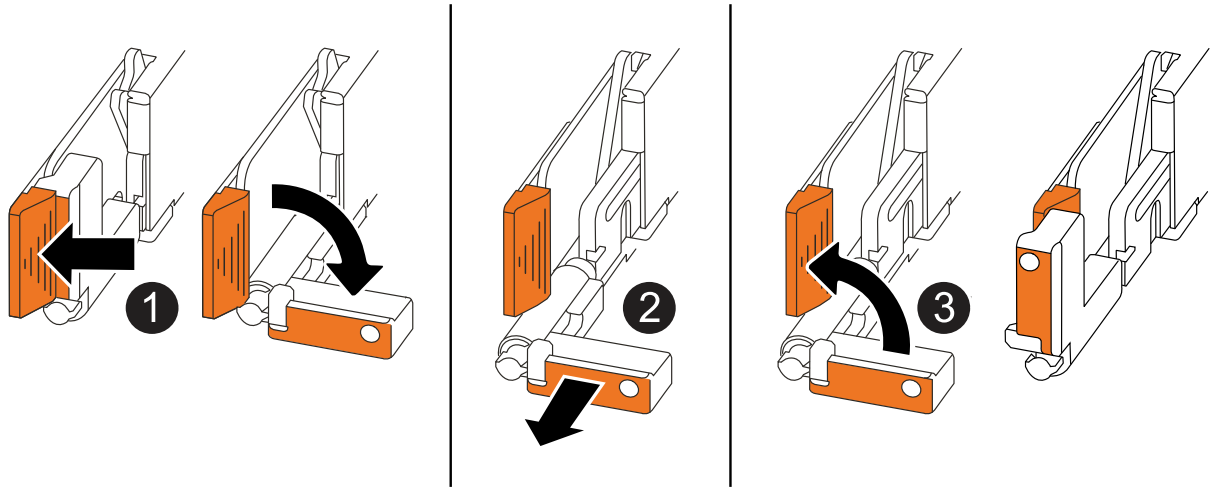

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Move the power supply

Move the power supply (PSU) to the replacement controller.

1. Move the PSU from the impaired controller:

   Make sure the left side controller handle is in the upright position to allow you access to the PSU.

Option 1: Move an AC PSU

To move an AC PSU, complete the following steps.

Steps

1. Remove the AC PSU from the impaired controller:

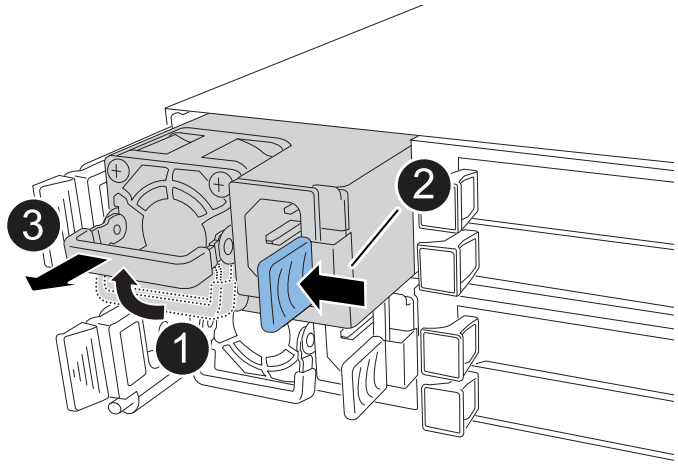

1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

2. Insert the PSU into the replacement controller:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.

Option 2: Move a DC PSU

To move a DC PSU, complete the following steps.

Steps

1. Remove the DC PSU from the impaired controller:

a. Rotate the handle up, to its horizontal position, and then grasp it.

b. With your thumb, press the terracotta tab to release the locking mechanism.

c. Pull the PSU out of the controller while using your other hand to support its weight.

> The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.

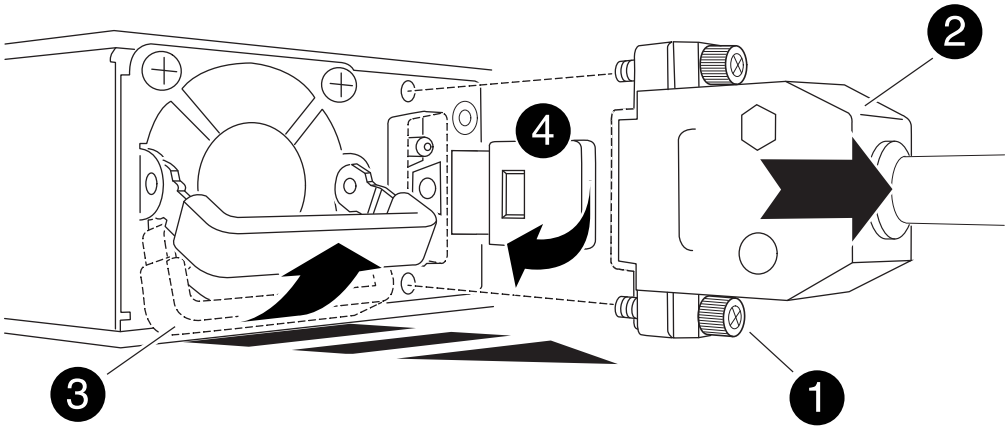

1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

2. Insert the PSU into the replacement controller:

a. Using both hands, support and align the edges of the PSU with the opening in the controller.

b. Gently slide the PSU into the controller until the locking tab clicks into place.

A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

> To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

c. Rotate the handle down, so it is out of the way of normal operations.

Step 3: Move the fans

Move the fans to the replacement controller.

1. Remove one of the fans from the impaired controller:

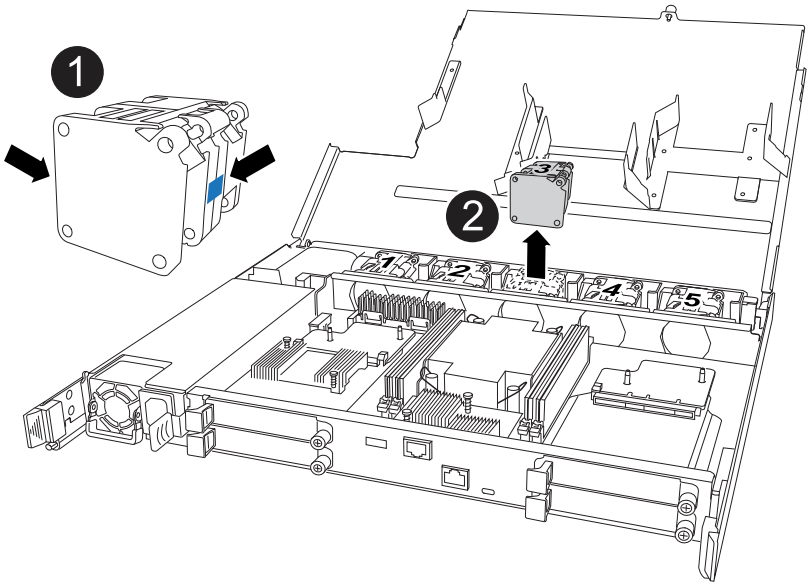

1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

2. Insert the fan into the replacement controller by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.
3. Repeat these steps for the remaining fans.

Step 4: Move the NV battery

Move the NV battery to the replacement controller.

1. Remove the NV battery from the impaired controller:

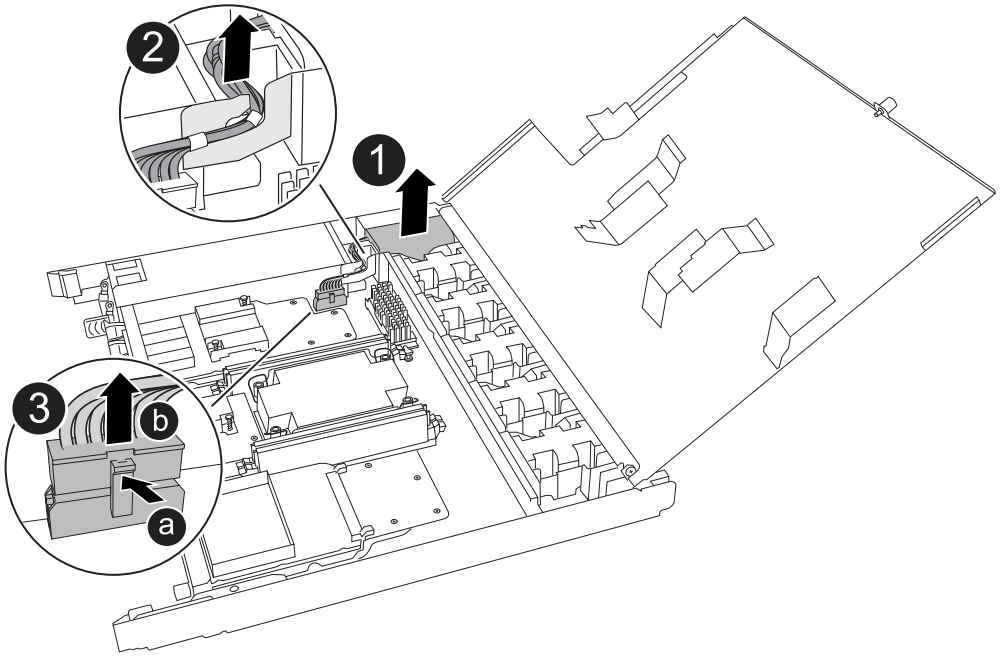

1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

2. Install the NV battery into the replacement controller:
   a. Plug the wiring connector into its socket.
   b. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   c. Place the NV battery into the compartment.

      The NV battery should sit flush in its compartment.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller.

If you have DIMM blanks, you do not need to move them, the replacement controller should come with them installed.

1. Remove one of the DIMMs from the impaired controller:

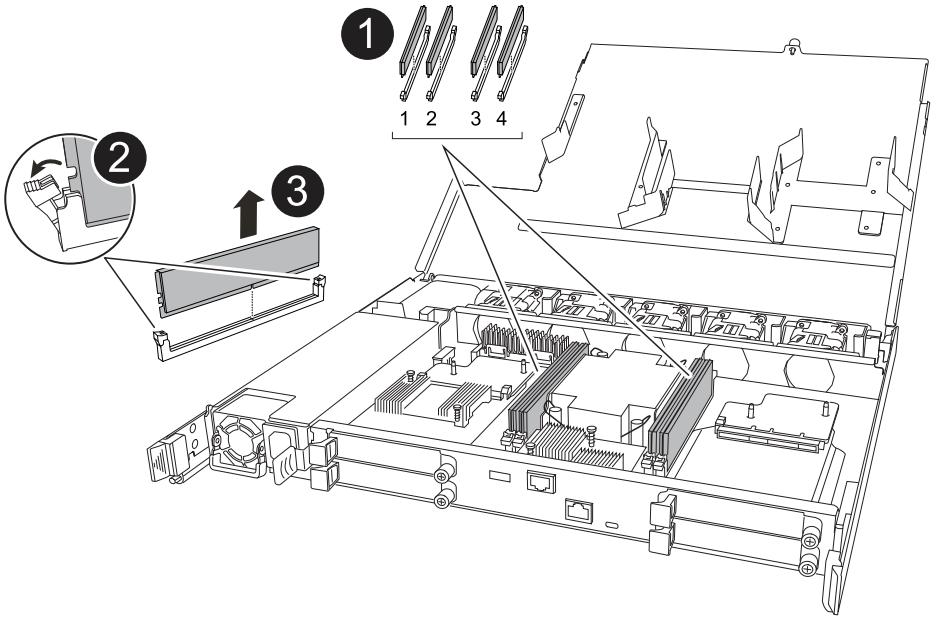

1	DIMM slot numbering and positions. Depending on your storage system model, you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller in the proper orientation. • Eject the DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

2. Install the DIMM in the replacement controller:
   a. Make sure that the DIMM ejector tabs on the connector are in the open position.
   b. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

      When inserted correctly, the DIMM goes in easily but fits tightly in the slot. If not, reinsert the DIMM.

   c. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
   d. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.
3. Repeat these steps for the remaining DIMMs.

Step 6: Move the boot media

Move the boot media to the replacement controller.

1. Remove the boot media from the impaired controller:

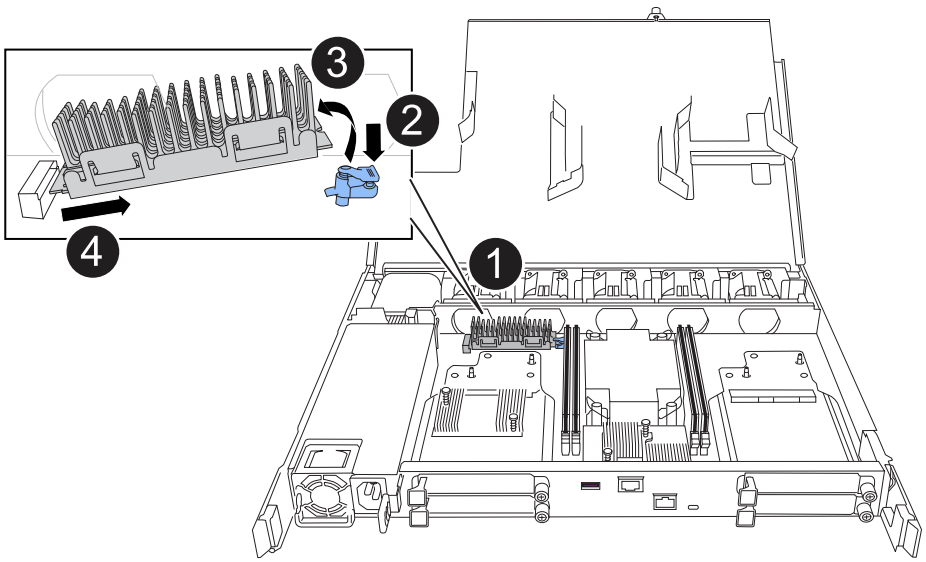

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

2. Install the boot media into the replacement controller:
   a. Slide the socket end of the boot media into its socket.
   b. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 7: Move the I/O modules

Move the I/O modules and any I/O blanking modules to the replacement controller.

1. Unplug cabling from one of the I/O modules.

   Make sure to label the cables so that you know where they came from.

2. Remove the I/O module from the impaired controller:

   Make sure that you keep track of which slot the I/O module was in.

   If you are removing the I/O module in slot 4, make sure the right side controller handle is in the upright position to allow you access to the I/O module.

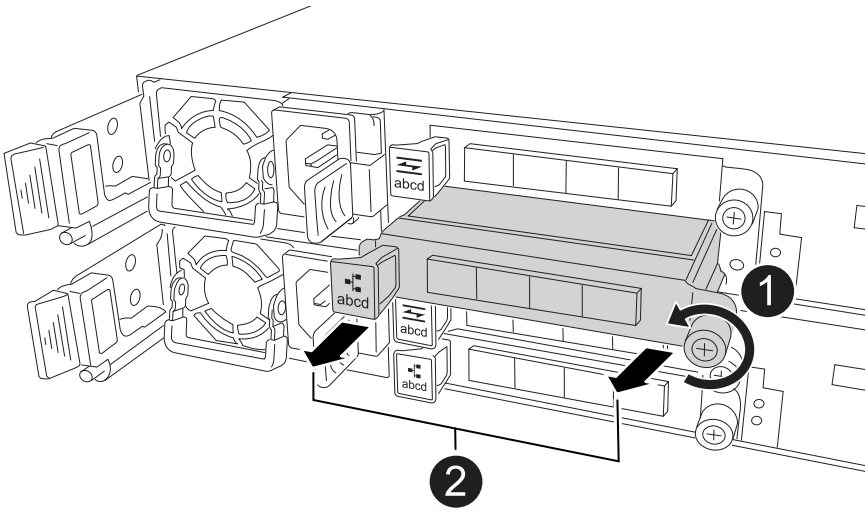

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

3. Install the I/O module into the replacement controller:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
4. Repeat these steps to move the remaining I/O modules and any I/O blanking modules to the replacement controller.

Step 8: Install the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

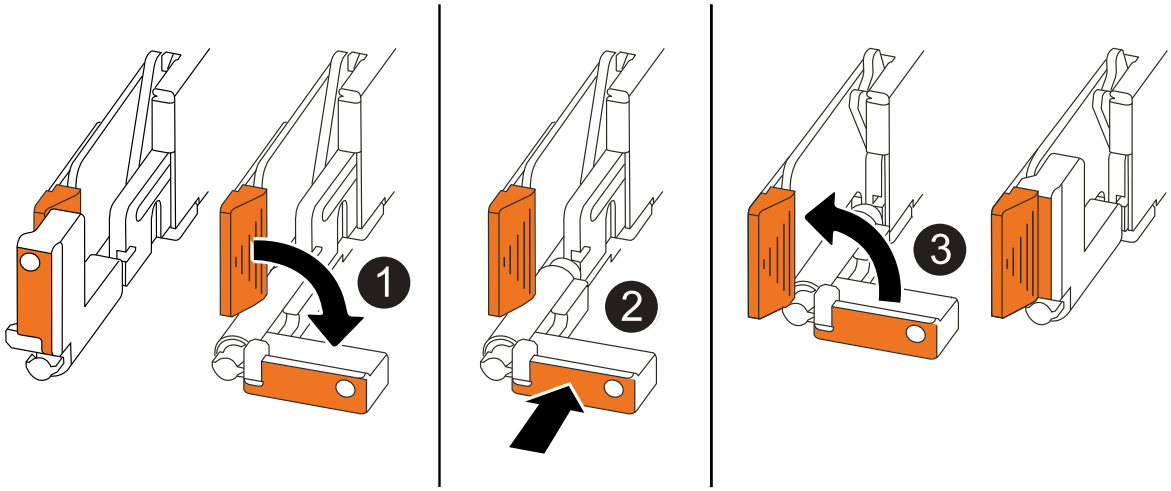

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Take the controller to the LOADER prompt by pressing CTRL-C to abort AUTOBOOT.
6. Set the time and date on the controller:

   Make sure you are at the controller's LOADER prompt.

   a. Display the date and time on the controller:

      ```
      show date
      ```

      Time and date default is in GMT. You have the option to display in local time and in 24hr mode.

   b. Set the current time in GMT:

      ```
      set time hh:mm:ss
      ```

      You can get the current GMT from the healthy node:

      ```
      date -u
      ```

   c. Set the current date in GMT:

      ```
      set date mm/dd/yyyy
      ```

      You can get the current GMT from the healthy node:
      ```
      date -u
      ```
7. Recable the controller as needed.
8. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

After you've replaced the impaired controller, you need to restore the system configuration.

Restore and verify the system configuration - ASA C30

Verify that the controller's HA configuration is active and functioning correctly in your ASA C30 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

You must verify the `HA` state of the controller and, if necessary, update the state to match your storage system configuration.

1. Boot to maintenance mode:

   ```
   boot_ontap maint
   ```

   a. Enter `y` when you see *Continue with boot?*.

      If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

5. If the displayed system state of the controller does not match your storage system configuration, set the `HA` state for the controller:

   ```
   ha-config modify controller ha
   ```

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed:

   ```
   ha-config show
   ```

Step 2: Verify disk list

1. Verify that the adapter lists the paths to all disks:

   ```
   storage show disk -p
   ```

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode:

```
halt
```

What's next?

After you've restored and verified your system configuration, you need to give back the controller.

Give back the controller - ASA C30

Return control of storage resources to the replacement controller so your ASA C30 storage system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption, Onboard Key Manager (OKM) encryption, or External Key Manager (EKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`

   If you encounter errors, contact NetApp Support.

9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

c. Verify status of the keys: `security key-manager key query -restored false`

   The output should show no results when when properly synchronized.

   If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

External key manager (EKM)

Reset encryption and return the controller to normal operation.

Steps

1. If the root volume is encrypted with External Key Manager and the console cable is connected to the replacement node, enter `boot_ontap menu` and select option `11`.
2. If these questions appear, answer `y` or `n` as appropriate:

   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}

   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}

   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}

   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}

   Do you know the KMIP server address? {y/n}

   Do you know the KMIP port? {y/n}

   Contact NetApp Support if you have issues.

3. Supply the information for:
   - The client certificate (client.crt) file contents
   - The client key (client.key) file contents
   - The KMIP server CA(s) (CA.pem) file contents
   - The IP address for the KMIP server
   - The port for the KMIP server
4. Once the system processes, you see the Boot Menu. Select '1' for normal boot.
5. Check the takeover status: `storage failover show`
6. Ensure any core dumps on the repaired node are saved by going to advanced mode `set -privilege advanced` and then `run local partner nosavecore`.
7. Return the impaired controller to normal operation by giving back its storage: `storage failover`

```
giveback -ofnode impaired_node_name
```

8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
9. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

After you've transferred the ownership of storage resources to the replacement controller, you need to complete the controller replacement procedure.

Complete controller replacement - ASA C30

To complete the controller replacement for your ASA C30 storage system, first restore the NetApp Storage Encryption configuration (if necessary) and install the required licenses on the new controller. Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, register the new controller's serial number and then return the failed part to NetApp.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

  Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The licenses keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.
- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs, register the serial number, and check cluster health

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - ASA C30

Replace a DIMM in your ASA C30 storage system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

Before you begin

- Ensure all other components in the storage system are working correctly; if not, contact NetApp Support before continuing."
- You must replace the failed FRU component with a replacement FRU component you received from your provider.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

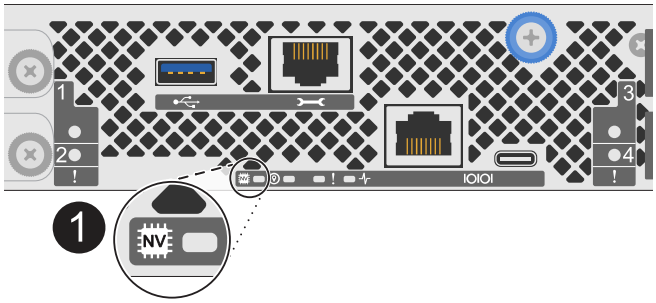

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

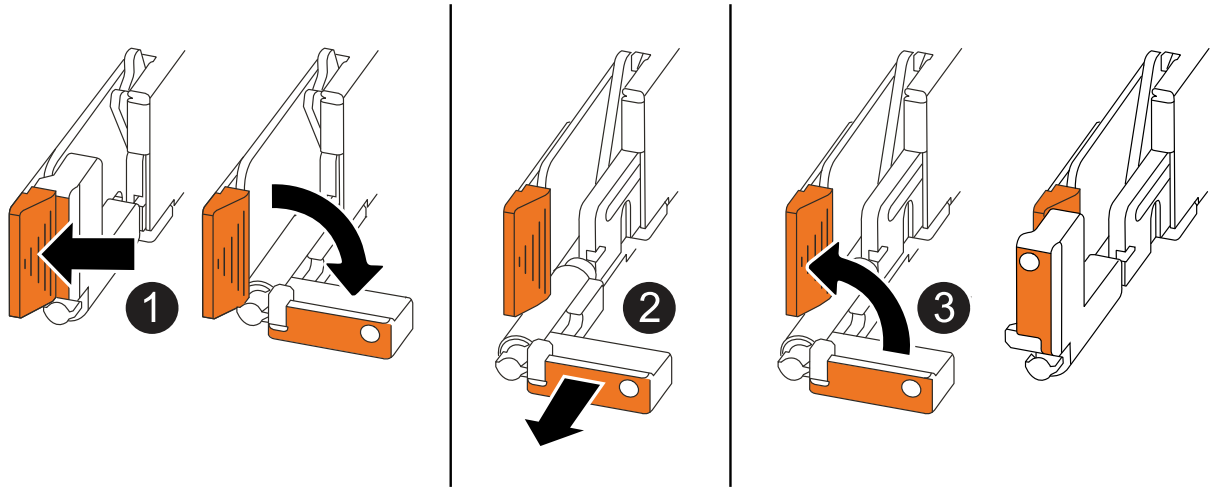

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace a DIMM

To replace a DIMM, locate the faulty DIMM inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller and identify the faulty DIMM.

> Consult either the Netapp Hardware Universe or the FRU map on the cover of the controller for exact DIMM locations.

3. Remove the faulty DIMM:

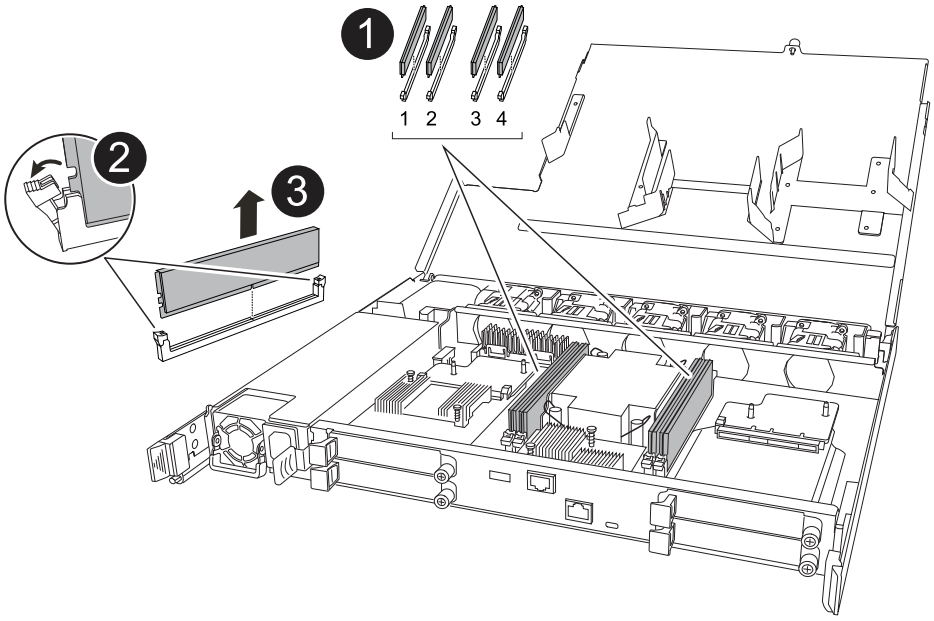

1	DIMM slot numbering and positions. Depending on your storage system model you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM using the same orientation. • Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

4. Install the replacement DIMM:
    a. Remove the replacement DIMM from its antistatic shipping bag.
    b. Make sure that the DIMM ejector tabs on the connector are in the open position.
    c. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

        The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

        When inserted correctly, the DIMM goes in easily but fits tightly in the slot. Reinsert the DIMM if you feel it is not inserted correctly.

    d. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
    e. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

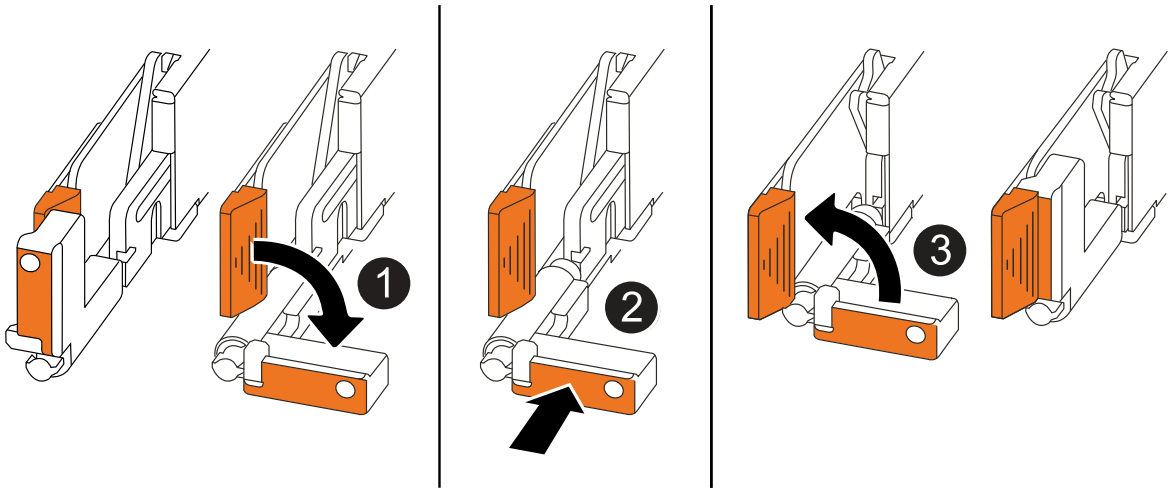

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

> Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

> Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

> Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

> The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a drive - ASA C30

Replace a drive in your ASA C30 storage system when a drive fails or requires an upgrade. The replacement process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed drive nondisruptively while I/O is in progress.

Before you begin

- The drive that you are installing must be supported by your storage system.

  NetApp Hardware Universe

- If self-encrypting drive (SED) authentication is enabled, you must use the SED replacement instructions in the ONTAP documentation.

  Instructions in the ONTAP documentation describe additional steps you must perform before and after replacing an SED.

  NetApp encryption overview with the CLI

- All other components in the storage system must be functioning properly, if not, you must contact NetApp Support before continuing with this procedure.
- Verify that the drive you are removing is failed.

  You can verify that the drive is failed by running the `storage disk show -broken` command. The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  Depending on the drive type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

About this task

- When replacing a failed drive, you must wait 70 seconds between the removal of the drive and the insertion of the replacement drive to allow the storage system to recognize that a drive was removed.
- The best practice is to have the current version of the Disk Qualification Package (DQP) installed before hot-swapping a drive.

  Having the current version of the DQP installed allows your system to recognize and use newly qualified drives. This avoids system event messages about having noncurrent drive information and prevention of drive partitioning because drives are not recognized. The DQP also notifies you of noncurrent drive firmware.

  NetApp Downloads: Disk Qualification Package

- The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.

NetApp Downloads: Disk Shelf Firmware

NetApp Downloads: Disk Drive Firmware

> Do not revert firmware to a version that does not support your shelf and its components.

- Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

  > Drive firmware checks occur every two minutes.

- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Steps

1. Properly ground yourself.
2. Remove the bezel from the front of the storage system.
3. Physically identify the failed drive.
   - When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the shelf operator display panel and the failed drive illuminate.
   - The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:

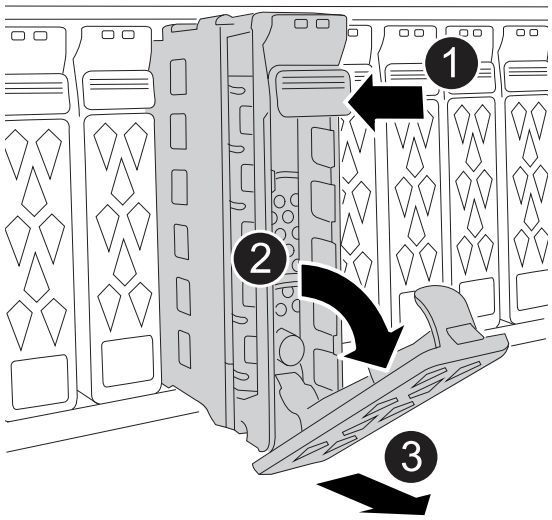

1	Press the release button on the drive face to open the cam handle.

2	Rotate the cam handle downward to disengage the drive from the midplane.
3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Gently push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

7. Verify that the drive's activity (green) LED is illuminated.

   When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat Steps 3 through 7.
9. Reinstall the bezel on the front of the storage system.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support NetApp Support if you need the RMA number or additional help with the replacement procedure.

Replace a fan module - ASA C30

Replace a fan module in your ASA C30 storage system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs

are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

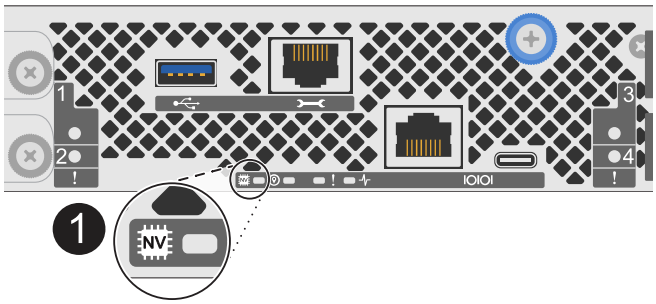

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

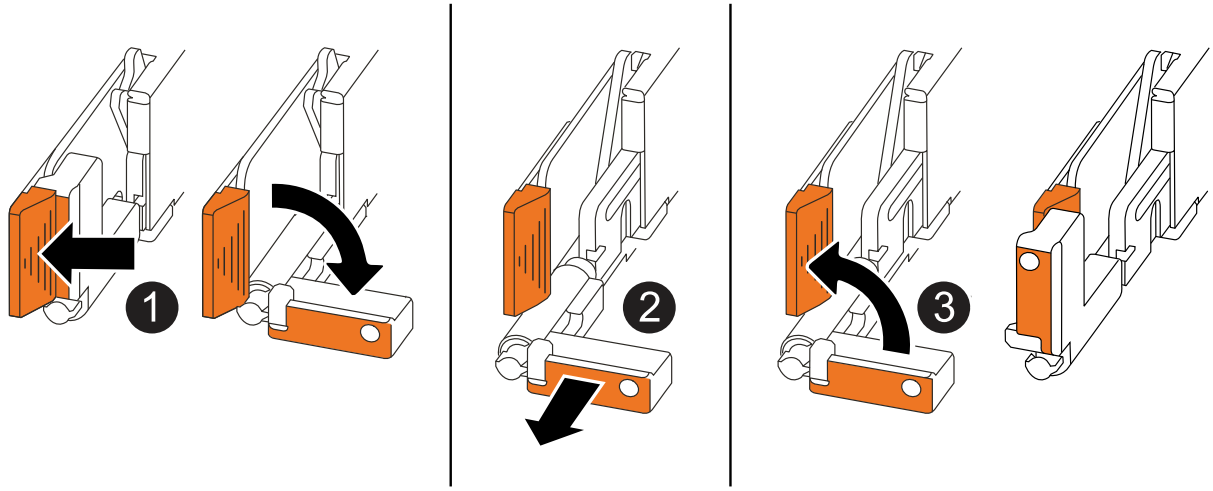

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace fan

To replace a fan, remove the failed fan and replace it with a new fan.

Steps

1. Identify the fan that you must replace by checking the console error messages.
2. Remove the failed fan:

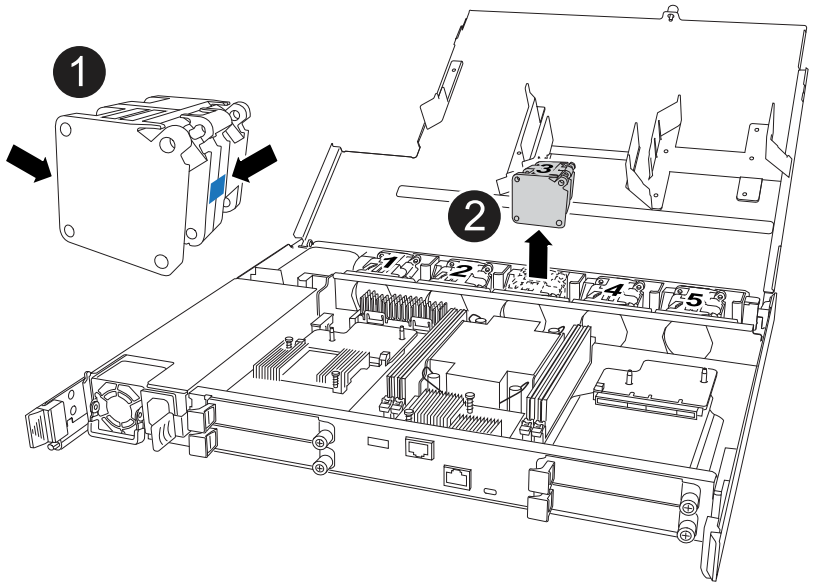

1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

3. Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

Step 4: Reinstall the controller module

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

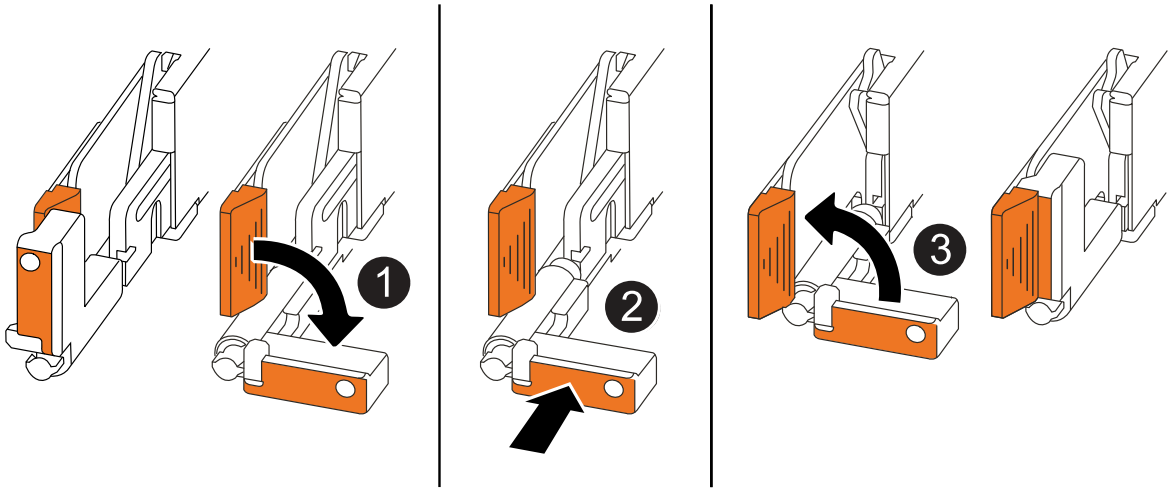

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```

8. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Overview of I/O module maintenance - ASA C30

The ASA C30 storage system offers flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding, hot-swapping, or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your storage system with the same type of I/O module, or with a different type of I/O module. You can hot-swap a cluster and HA I/O module when your storage system meets specific requirements. You can also add an I/O module to a storage system with available slots.

- Add an I/O module

  Adding additional I/O modules can improve redundancy, helping to ensure that the storage system remains operational even if one I/O module fails.

- Hot-swap a cluster and HA I/O module

  Hot-swapping a failed cluster and HA I/O module can restore the storage system to its optimal operating state. Hot-swapping is done without having to manually take over the impaired controller.

  To use this procedure, your storage system must be running ONTAP 9.17.1 or later and meet specific system requirements.

- Replace an I/O module

  Replacing a failing I/O module can restore the storage system to its optimal operating state.

Add an I/O module - ASA C30

Add an I/O module to your ASA C30 storage system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your ASA C30 storage system if there are available slots. If all slots are fully

populated, you can replace an existing module to add a new one.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, remove the I/O blanking module from the target slot.

   Unused I/O slots should have blanking module installed to prevent possible thermal issues and for EMC compliance.

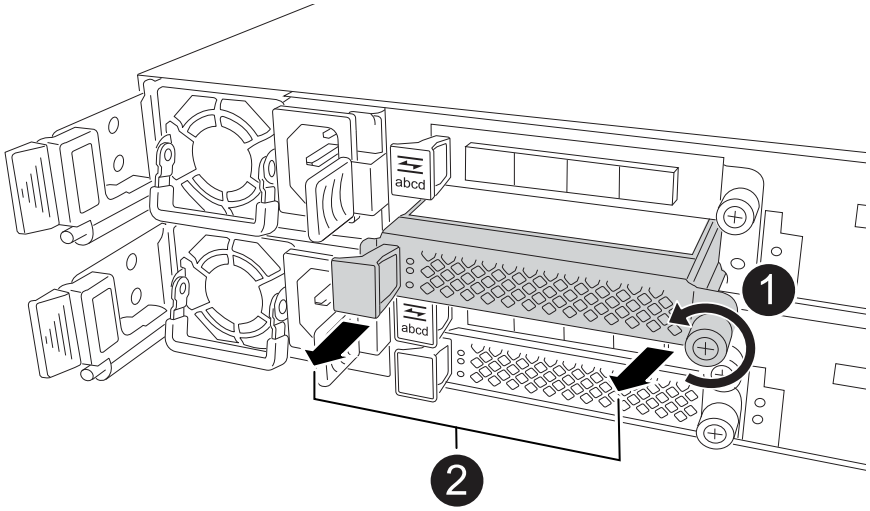

1	On the I/O blanking module, turn the thumbscrew counterclockwise to loosen.
2	Pull the I/O blanking module out of the controller using the tab on the left and the thumbscrew.

3. Install the new I/O module:
   a. Align the I/O module with the edges of the controller slot opening.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O module.

   c. Turn the thumbscrew clockwise to tighten.
4. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.

5. Reboot the impaired controller from the LOADER prompt: `bye`

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

6. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name.
```

7. Repeat these steps to add an I/O module to the other controller.
8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation: +

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, unplug any cabling on the target I/O module.
3. Remove the target I/O module from the controller:

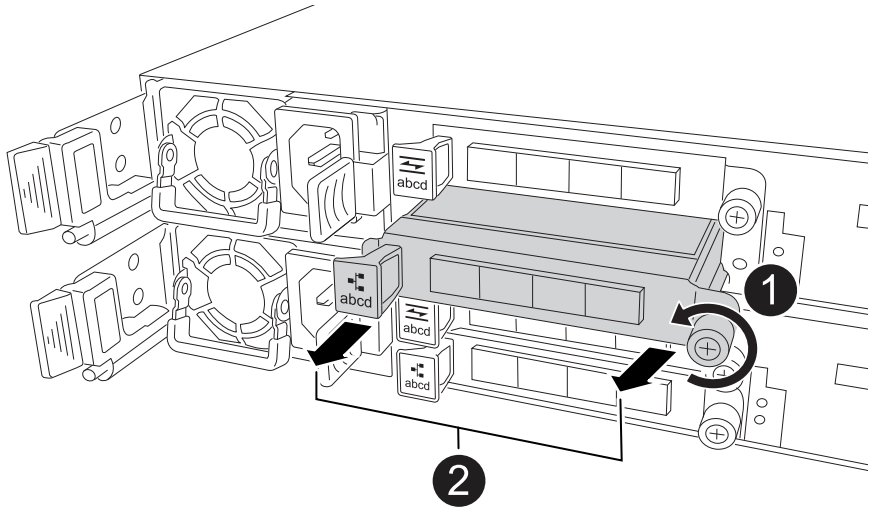

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the new I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
5. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
6. Repeat the I/O module remove and install steps to add any additional I/O modules in the controller.
7. Reboot the impaired controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.
8. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```
9. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```
10. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```
11. If you installed a NIC module, specify the usage mode for each port as *network*:

    ```
    storage port modify -node node_name -port port_name -mode network
    ```
12. Repeat these steps for the other controller.

Hot-swap the I/O module used for cluster and HA traffic - ASA C30

The cluster and HA I/O module supports interconnects for clustering and high-availability. You can hot-swap the module in your ASA C30 storage system when the module fails and if your storage system meets specific requirements.

To hot-swap a module, you ensure your storage system meets the procedure requirements, prepare the storage system and I/O module in slot 4, hot-swap the failed module for an equivalent one, bring the replacement module online, restore the storage system to normal operation, and return the failed module to NetApp.

About this task

- Hot-swapping the cluster and HA I/O module means that you do not have to perform a manual takeover; the impaired controller (the controller with the failed cluster and HA I/O module) has automatically taken over the healthy controller.

  When the impaired controller has taken over the healthy controller, the only way to recover without an outage is to hot-swap the module.

- It is critical to apply the commands to the correct controller when you are hot-swapping the cluster and HA I/O module:
  - The *impaired controller* is the controller on which you are hot-swapping the cluster and HA I/O module and it is the controller that has taken over the healthy controller.
  - The *healthy controller* is the HA partner of the impaired controller and it is the controller that was taken over by the impaired controller.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, make sure your storage system meets all requirements.

> If your storage system does not meet all requirements, you must use the replace an I/O module procedure.

- Your storage system must be running ONTAP 9.17.1 or later.
- The I/O module that failed must be a cluster and HA I/O module in slot 4 and you must be replacing it with an equivalent cluster and HA I/O module. You cannot change the I/O module type.
- Your storage system configuration must have only one cluster and HA I/O module located in slot 4, not two cluster and HA I/O modules.
- Your storage system must be a two-node (switchless or switched) cluster configuration.
- The controller with the failed cluster and HA I/O module (the impaired controller) must have already taken over the healthy partner controller. The takeover should have occurred automatically if the I/O module is failed.

  For two-node clusters, the storage system cannot discern which controller has the failed I/O module, so either controller might initiate the takeover. The cluster and HA I/O module hot-swap procedure is only supported when the controller with the failed I/O module (the impaired controller) has taken over the healthy controller.

You can verify that the impaired controller successfully took over the healthy controller by entering the `storage failover show` command.

If you are not sure which controller has the failed I/O module, contact NetApp Support.

- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Step 2: Prepare the storage system and I/O module slot 4

Prepare the storage system and I/O module slot 4 so that it is safe to remove the failed cluster and HA I/O module:

Steps

1. Properly ground yourself.
2. Unplug cabling from the failed cluster and HA I/O module.

   Make sure to label the cables so that later in this procedure you can reconnect them to the same ports.

3. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<number of hours down>h
   ```

   For example, the following AutoSupport message suppresses automatic case creation for two hours:

   ```
   node2::> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

4. Disable automatic giveback:
   a. Enter the following command from the console of the impaired controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
5. Prepare the failed cluster and HA module in slot 4 for removal by removing it from service and powering it off:
   a. Enter the following command:

      ```
      system controller slot module remove -node impaired_node_name -slot slot_number
      ```

   b. Enter `y` when you see the prompt *Do you want to continue?*

      For example, the following command prepares the module in slot 4 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

```
node2::> system controller slot module remove -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered off for removal.

Do you want to continue? {y|n}: y

The module has been successfully removed from service and powered
off. It can now be safely removed.
```

6. Verify the failed cluster and HA module in slot 4 is powered off:

   `system controller slot module show`

   The output should show *`powered-off`* in the status column for the failed module in slot 4.

Step 3: Replace the failed cluster and HA I/O module

Replace the failed cluster and HA I/O module in slot 4 with an equivalent I/O module:

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the failed cluster and HA I/O module from the impaired controller:

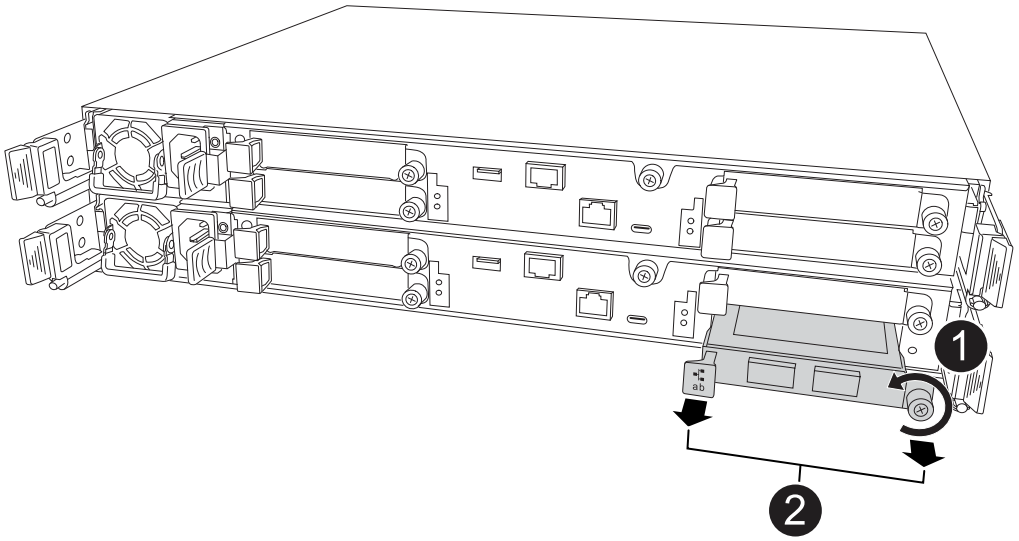

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew on the right.

3. Install the replacement cluster and HA I/O module into slot 4:

a. Align the I/O module with the edges of the slot.
b. Gently push the I/O module all the way into the slot, making sure to properly seat the I/O module into the connector.

   You can use the tab on the left and the thumbscrew on the right to push in the I/O module.

c. Turn the thumbscrew clockwise to tighten.

4. Cable the cluster and HA I/O module.

Step 4: Bring the replacement cluster and HA I/O module online

Bring the replacement cluster and HA I/O module in slot 4 online, verify the module ports initialized successfully, verify slot 4 is powered on, and then verify the module is online and recognized.

Steps

1. Bring the replacement cluster and HA I/O module online:
   a. Enter the following command:

      `system controller slot module insert -node impaired_node_name -slot slot_name`

   b. Enter `y` when you see the prompt, *Do you want to continue?*

      The output should confirm the cluster and HA I/O module was successfully brought online (powered on, initialized, and placed into service).

      For example, the following command brings slot 4 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered on and initialized.

Do you want to continue? {y|n}: `y`

The module has been successfully powered on, initialized and placed
into service.
```

2. Verify that each port on the cluster and HA I/O module successfully initialized:

   `event log show -event *hotplug.init*`

   > It might take several minutes to allow for any required firmware updates and port initialization.

   The output should show a hotplug.init.success EMS event logged for each port on the cluster and HA I/O module with *`hotplug.init.success:`* in the *`Event`* column.

For example, the following output shows initialization succeeded for cluster and HA I/O module ports e4b and e4a:

```
node2::> event log show -event *hotplug.init*

Time                Node             Severity      Event

------------------- ---------------- -------------
---------------------------

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4b" in slot 4 succeeded

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4a" in slot 4 succeeded

2 entries were displayed.
```

3. Verify I/O module slot 4 is powered on and ready for operation:

   `system controller slot module show`

   The output should show slot 4 status as *`powered-on`* and therefore ready for operation of the replacement cluster and HA I/O module.

4. Verify that the replacement cluster and HA I/O module is online and recognized.

   Enter the command from the console of the impaired controller:

   `system controller config show -node local -slot4`

   If the replacement cluster and HA I/O module was successfully brought online and is recognize, the output shows I/O module information, including port information, for slot 4.

   For example, you should see output similar to the following:

```
node2::> system controller config show -node local -slot 4

Node: node2
Sub- Device/
Slot slot Information
---- ---- ------------------------------
   4    - Dual 40G/100G Ethernet Controller CX6-DX
                  e4a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2807GJFM-B
                  e4b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2809G26F-A
                  Device Type:        CX6-DX PSID(NAP0000000027)
                  Firmware Version:   22.44.1700
                  Part Number:        111-05341
                  Hardware Revision:  20
                  Serial Number:      032403001370
```

Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the healthy controller, restoring automatic giveback, and reenabling AutoSupport automatic case creation.

Steps

1. Return the healthy controller (the controller that was taken over) to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode healthy_node_name
   ```

2. Restore automatic giveback from the console of the impaired controller (the controller that took over the healthy controller):

   ```
   storage failover modify -node local -auto-giveback true
   ```

3. If AutoSupport is enabled, restore automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=end
   ```

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an I/O module in your ASA C30 storage system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

Use this procedure to replace a failed I/O module.

Before you begin

All other components in the storage system must be functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback false
```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace a failed I/O module

To replace a failed I/O module, locate it in the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug cabling from the failed I/O module.

   Make sure to label the cables so that you know where they came from.

3. Remove the failed I/O module from the controller:

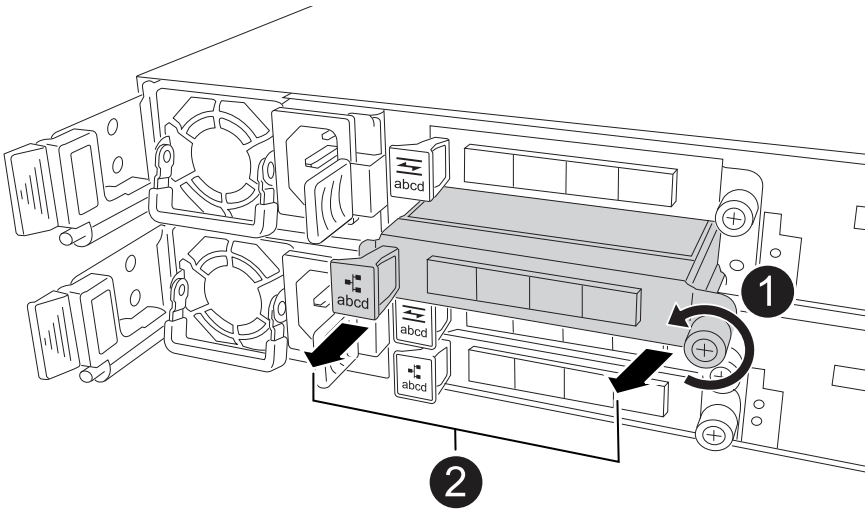

1	Turn the I/O module thumbscrew counterclockwise to loosen.

2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the replacement I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
5. Cable the I/O module.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt: `bye`

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the node to normal operation: `storage failover giveback -ofnode` *`impaired_node_name`*
3. Restore automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback` *`true`*

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NV battery - ASA C30

Replace the NV battery in your ASA C30 storage system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

To replace the NV battery, you must remove the controller, remove the faulty battery, install the replacement battery, and then reinstall the controller.

Before you begin

All other components in the storage system must be functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected

storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

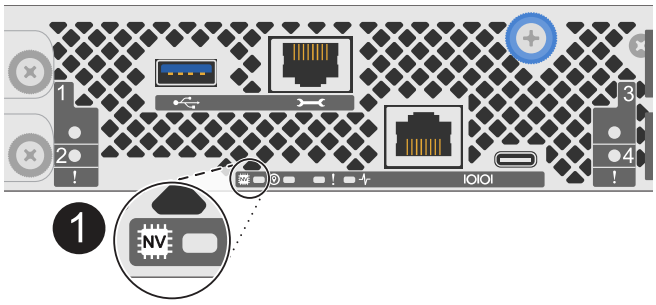

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

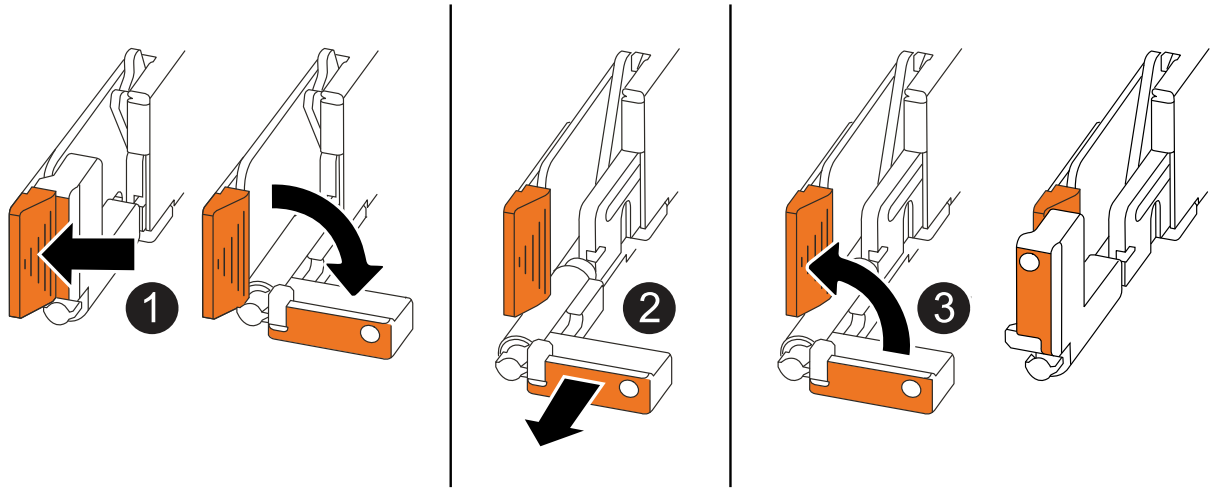

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the NV battery

Remove the failed NV battery from the controller and install the replacement NV battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the NV battery.
3. Remove the NV battery:

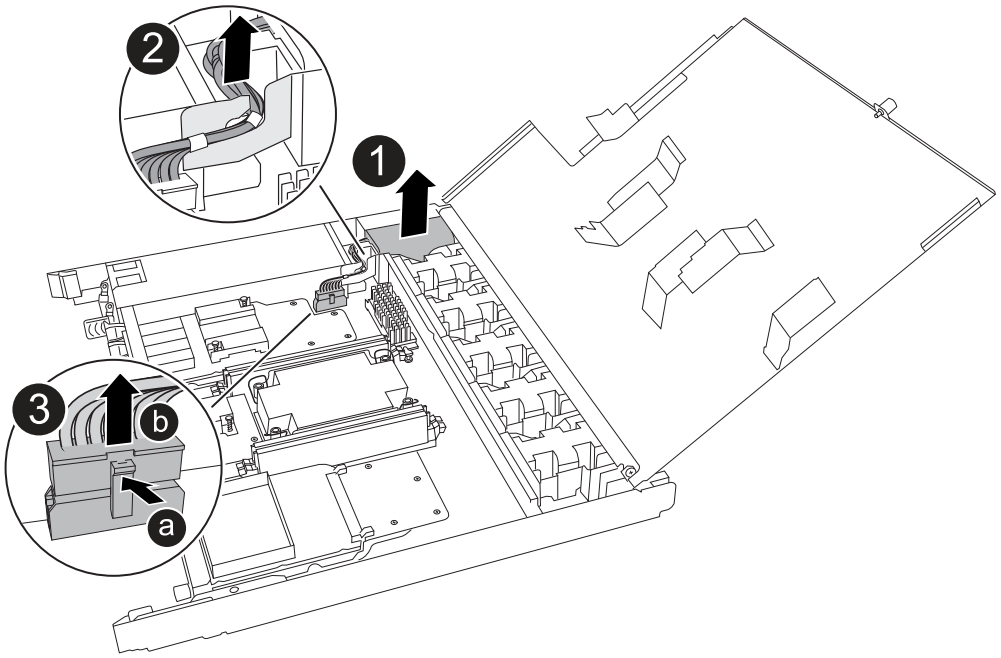

1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

4. Install the replacement NV battery:
   a. Remove the replacement battery from its package.
   b. Plug the wiring connector into its socket.
   c. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   d. Place the NV battery into its compartment.

      The NV battery should sit flush in its compartment.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

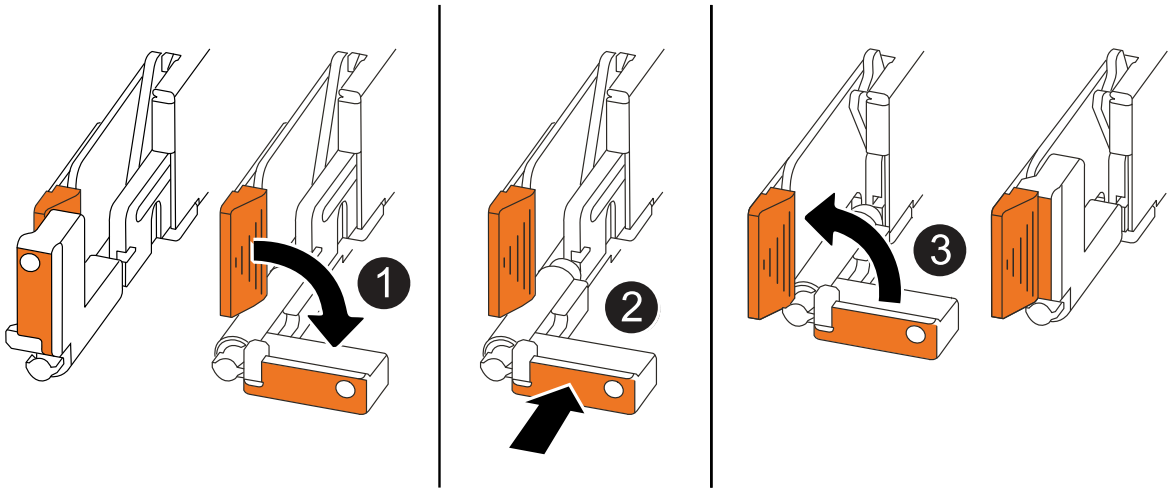

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.

The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```

8. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - ASA C30

Replace an AC or DC power supply unit (PSU) in your ASA C30 storage system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cord, replacing the faulty PSU, and then reconnecting it to the power source.

About this task

- This procedure is written for replacing one PSU at a time.

The PSUs are redundant and hot-swappable.

- IMPORTANT: Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.
- Use the appropriate procedure for your type of PSU: AC or DC.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the power cord from the PSU by opening the power cord retainer, and then unplug the power cord from the PSU.

   PSUs do not have a power switch.

4. Remove the PSU:

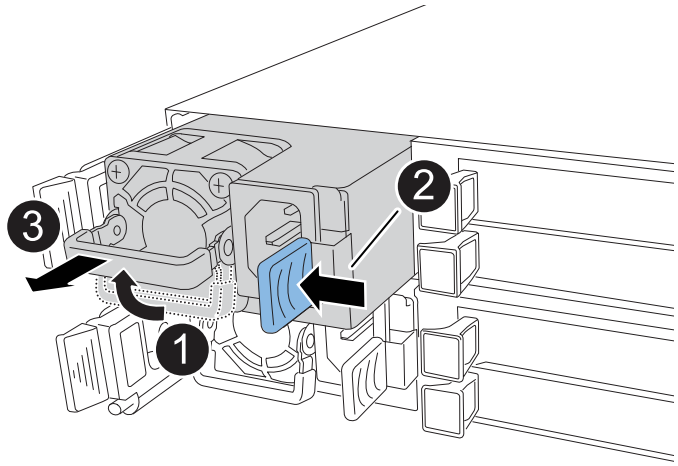

1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

5. Install the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

c. Rotate the handle down, so it is out of the way of normal operations.

6. Reconnect the power cord to the PSU and secure the power cord with the power cord retainer.

   Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the PSU:

   PSUs do not have a power switch.

   a. Unscrew the two thumb screws on the D-SUB DC power cord connector.

      The illustration and table in step 4 shows the two thumb screws (item #1) and the D-SUB DC power cord connector (item #2).

   b. Unplug the cord from the PSU and set it aside.

4. Remove the PSU:

   a. Rotate the handle up, to its horizontal position, and then grasp it.

   b. With your thumb, press the terracotta tab to release the locking mechanism.

   c. Pull the PSU out of the controller while using your other hand to support its weight.

   The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.

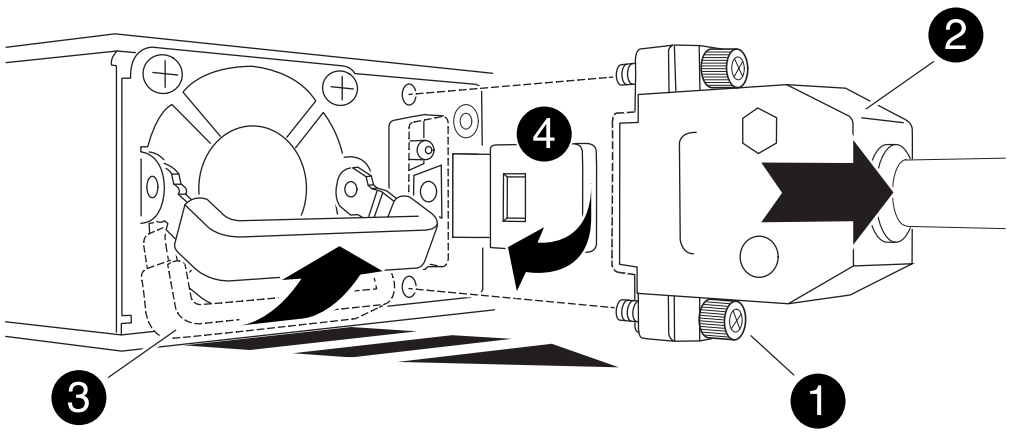

1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

5. Insert the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently slide the PSU into the controller until the locking tab clicks into place.

      A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.
6. Reconnect the D-SUB DC power cord:

   Once power is restored to the PSU, the status LED should be green.

   a. Plug the D-SUB DC power cord connector into the PSU.
   b. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - ASA C30

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your ASA C30 storage system to ensure that services and applications relying on accurate time synchronization remain operational.

You replace the real-time clock (RTC) battery in the controller so that your storage system's services and applications that depend on accurate time synchronization continue to function.

Before you begin

All other components in the storage system must be functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

About this task

- You can use this procedure with all versions of ONTAP supported by your storage system.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

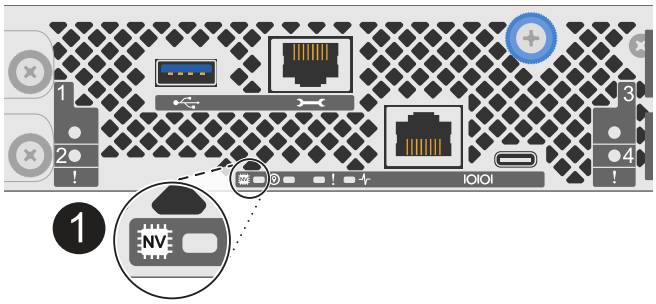

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

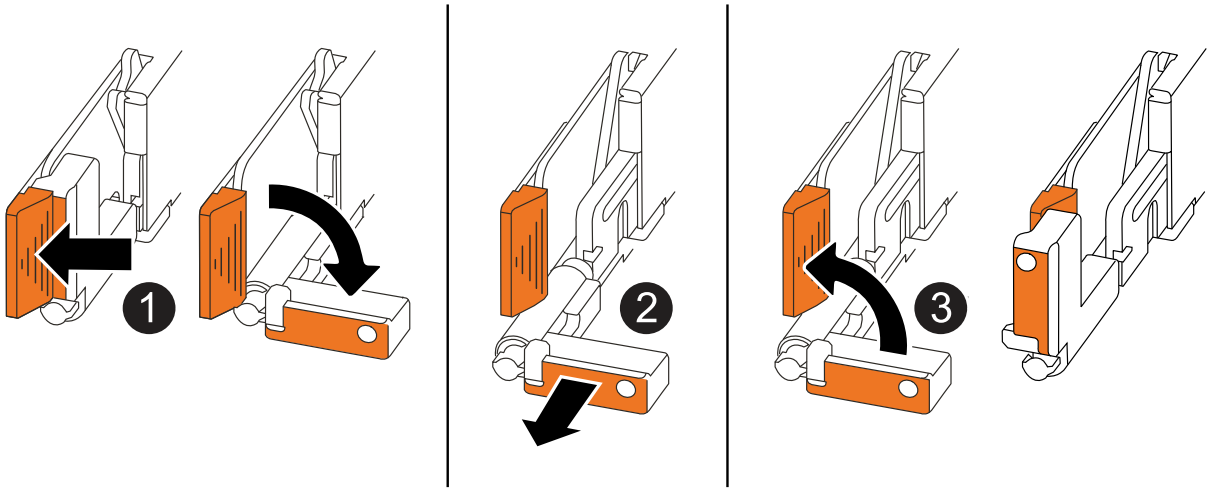

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the RTC battery

Remove the failed RTC battery and install the replacement RTC battery.

Steps

1. Locate the RTC battery.
2. Remove the RTC battery:

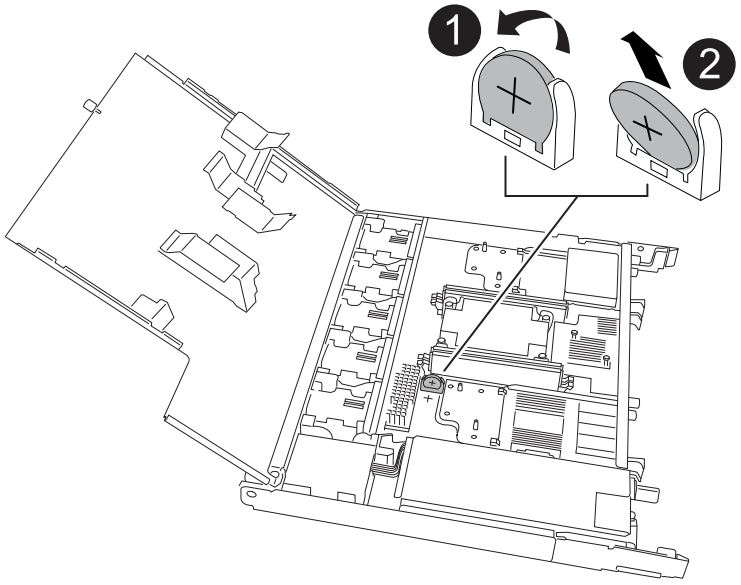

1	Gently rotate the RTC battery at an angle away from its holder.
2	Lift the RTC battery out of its holder.

3. Install the replacement RTC battery:
   a. Remove the replacement battery from the antistatic shipping bag.
   b. Position the battery so that the plus sign on the battery faces out to correspond with the plus sign on the motherboard.
   c. Insert the battery into the holder at an angle, and then push it into an upright position so it is fully seated in the holder.
   d. Visually inspect the battery to make sure that it is completely seated in its holder and that the polarity is correct.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

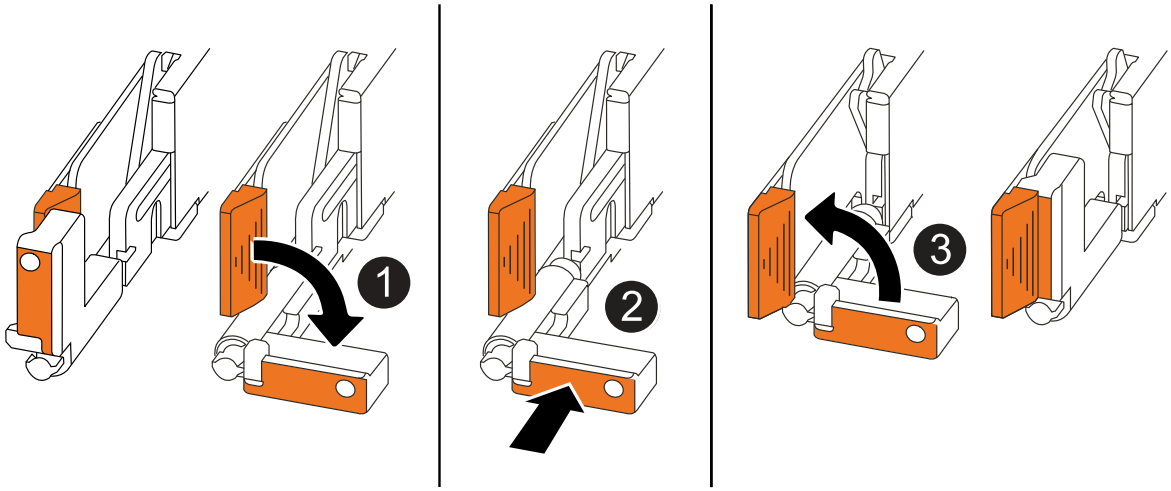

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode` *`impaired_node_name`*

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Reset the time and date on the controller

After replacing the RTC battery, inserting the controller and powering on first BIOS reset, you will see the following error messages:
`RTC date/time error. Reset date/time to default`
`RTC power failure error`
These messages are expected and you can continue with this procedure.

1. On the healthy controller, check the date and time with the `cluster date show` command.

   If your storage system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*.

2. On the impaired controller, at the LOADER prompt, check the time and date: `cluster date show`
   a. If necessary, modify the date: `set date` *`mm/dd/yyyy`*
   b. If necessary, set the time, in GMT: `set time hh:mm:ss`
   c. Confirm the date and time.
3. At the LOADER prompt, enter `bye` to reinitialize the I/O modules, other components, and let the controller reboot.
4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

FAS systems

FAS50 systems

Install and setup

Installation and setup workflow - FAS50

To install and set up your FAS50 storage system, you must review the installation requirements, prepare your site, install and cable the hardware components, power on the storage system, and set up the ONTAP cluster.

Review the installation requirements

Before installing your storage system, it must meet the installation requirements.

Prepare for installation

To prepare for installation, get the site ready, check environmental and electrical requirements, and ensure there's enough rack space. Then, unpack the equipment, compare contents to the packing slip, and register the hardware to access support benefits.

Install the hardware

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system and shelves in the cabinet or telco rack.

Cable the hardware

To cable the hardware, connect the controllers to your network and then to your shelves.

Power on the storage system

To power on your storage system, power on each shelf and assign a unique shelf ID as needed, and then power on the controllers.

Set up your cluster

After you've powered on your storage system, you set up your cluster.

Installation requirements - FAS50

Review the requirements for your FAS50 storage system.

Equipment needed for install

To install your storage system, you need the following equipment and tools.

- Access to a Web browser to configure your storage system
- Electrostatic discharge (ESD) strap
- Flashlight
- Laptop or console with a USB/serial connection
- Phillips #2 screwdriver

Lifting precautions

Storage systems and shelves are heavy. Exercise caution when lifting and moving these items.

Storage system weight

Take the necessary precautions when moving or lifting your storage system.

An A1K storage system can weigh up to 62.83 lbs (28.5 kg). To lift the storage system, use two people or a hydraulic lift.

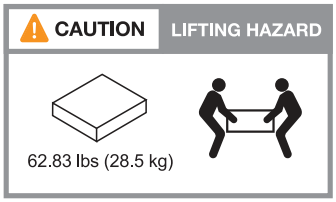

The storage system can weigh up to 53.8 lbs (24.4 kg). To lift the storage system, use two people or a hydraulic lift.

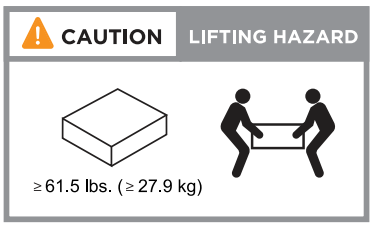

Shelf weight

Take the necessary precautions when moving or lifting your shelf.

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

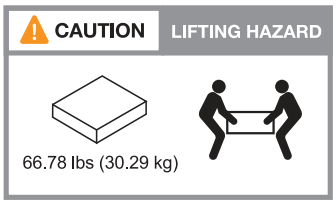

A DS460C shelf can weigh up to 260.4 lbs (181.1 kg). To lift the shelf, you might need up to five people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

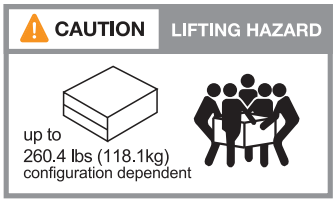

Related information

- Safety information and regulatory notices

What's next?

After you've reviewed the installation requirements and considerations for your storage system, you prepare for installation.

Prepare to install - FAS50

Prepare to install your FAS50 storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the storage system to access support benefits.

Step 1: Prepare the site

To install your storage system, ensure that the site and the cabinet or rack that you plan to use meet specifications for your configuration.

Steps

1. Use NetApp Hardware Universe to confirm that your site meets the environmental and electrical requirements for your storage system.
2. Make sure you have adequate cabinet or rack space for your storage system, shelves, and any switches:
3. Install any required network switches.

   See the Switch documentation for installation instructions and NetApp Hardware Universe for compatibility information.

Step 2: Unpack the boxes

After you've ensured that the site and the cabinet or rack that you plan to use for your storage system meet the required specifications, unpack all boxes and compare the contents to the items on the packing slip.

Steps

1. Carefully open all the boxes and lay out the contents in an organized manner.
2. Compare the contents you've unpacked with the list on the packing slip.

   You can get your packing list by scanning the QR code on the side of the shipping carton.

   The following items are some of the contents you might see in the boxes.

   Ensure that everything in the boxes matches the list on the packing slip. If there are any discrepancies, note them down for further action.

Hardware

- Bezel
- Storage system
- Rail kits with instructions (optional)
- Storage shelf (if you ordered additional storage)

Cables

- Management Ethernet cables (RJ-45 cables)
- Network cables
- Power cords
- Storage cables (if you ordered additional storage)
- USB-C serial console cable

Step 3: Register your storage system

After you've ensured that your site meets the requirements for your storage system specifications, and you've verified that you have all the parts you ordered, you should register your storage system.

Steps

1. Locate the System Serial Numbers (SSN) for every controller being installed. You can find the serial numbers in the following locations:
2. You can find the serial numbers in the following locations:
   - On the packing slip
   - In your confirmation email

3. Go to the NetApp Support Site.
4. Determine whether you need to register your storage system:

If you are a…	Follow these steps…
Existing NetApp customer	a. Sign in with your username and password. b. Select **Systems** > **My Systems**. c. Confirm that the new serial numbers are listed. d. If it is not, follow the instructions for new NetApp customers.
New NetApp customer	a. Click **Register Now**, and create an account. b. Select **Systems** > **Register Systems**. c. Enter the storage system's serial numbers and requested details. After your registration is approved, you can download any required software. The approval process might take up to 24 hours.

What's next?

After you've prepared to install your storage system, you install the hardware for your storage system.

Install the hardware - FAS50

After you prepare to install your FAS50 storage system, install the hardware for the storage system. First, install the rail kits. Then install and secure your storage system in a cabinet or telco rack.

Skip this step if your storage system came in a cabinet.

Before you begin

- Make sure you have the instructions packaged with the rail kit.
- Be aware of the safety concerns associated with the weight of the storage system and shelf.
- Understand that the airflow through the storage system enters from the front where the bezel or end caps are installed and exhausts out the rear where the ports are located.

Steps

1. Install the rail kits for your storage system and shelves as needed, using the instructions included with the kits.
2. Install and secure your storage system in the cabinet or telco rack:
   a. Position the storage system onto the rails in the middle of the cabinet or telco rack, and then support the storage system from the bottom and slide it into place.
   b. Make sure that the guiding pins of the cabinet or telco rack are securley in the chassis guide slots.
   c. Secure the storage system to the cabinet or telco rack using the included mounting screws.
3. Attach the bezel to the front of the storage system.
4. Install and secure the shelf as needed.
   a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

      If you are installing multiple shelves, place the first shelf directly above the controllers. Place the second shelf directly under the controllers. Repeat this pattern for any additional shelves.

   b. Secure the shelf to the cabinet or telco rack using the included mounting screws.

      What's next?

      After you've installed the hardware for your storage system, you cable the hardware.

Cable the hardware - FAS50

After you install your FAS50 storage system hardware, cable the controllers to the network and shelves.

Before you begin

Contact your network administrator for information about connecting the storage system to your network switches.

About this task

- The cabling graphics have arrow icons showing the proper orientation (up or down) of the cable connector pull-tab when inserting a connector into a port.

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it over and try again.

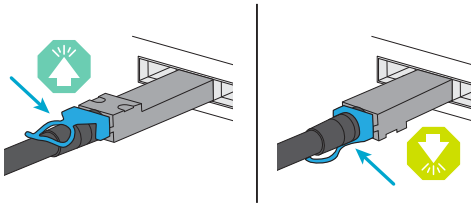

- If cabling to an optical switch, insert the optical transceiver into the controller port before cabling to the switch port.

Step 1: Cable the cluster/HA connections

Create the ONTAP cluster connections. For switchless clusters, connect the controllers to each other. For switched clusters, connect the controllers to the cluster network switches.

The cluster/HA cabling examples show common configurations.

If you do not see your configuration here, go to NetApp Hardware Universe for comprehensive configuration and slot priority information to cable your storage system.

Switchless cluster cabling

FAS50 with one 2-port 40/100 GbE I/O module

Steps

1. Cable the Cluster/HA interconnect connections:

   The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

   a. Cable controller A port e4a to controller B port e4a.

   b. Cable controller A port e4b to controller B port e4b.

100 GbE Cluster/HA interconnect cables

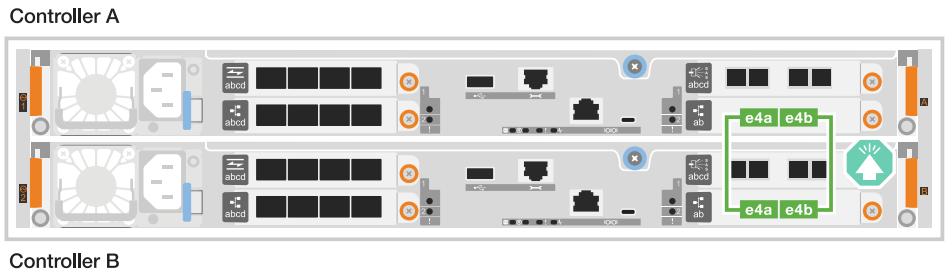

Switched cluster cabling

FAS50 with one 2-port 40/100 GbE I/O module

1. Cable the controllers to the cluster network switches:

   The cluster interconnect traffic and the HA traffic share the same physical ports(on the I/O module in slot 4). The ports are 40/100 GbE.

   a. Cable controller A port e4a to cluster network switch A.
   b. Cable controller A port e4b to cluster network switch B.
   c. Cable controller B port e4a to cluster network switch A.
   d. Cable controller B port e4b to cluster network switch B.

40/100 GbE Cluster/HA interconnect cables

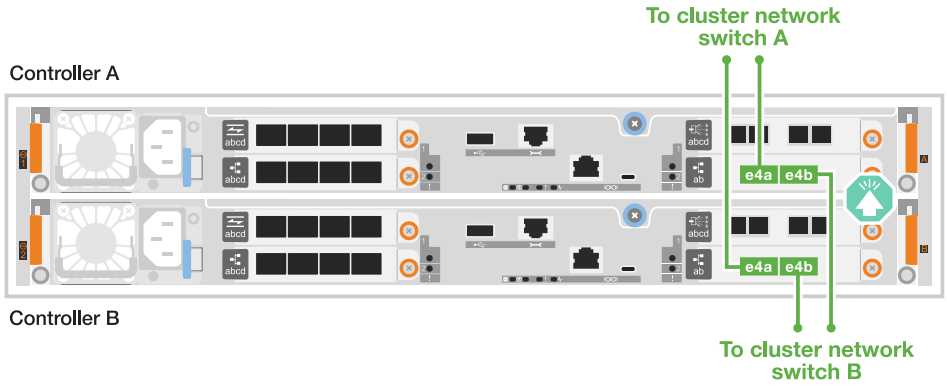

Step 2: Cable the host network connections

Cable the controllers to your Ethernet or FC host network.

The host network cabling examples show common configurations.

If you do not see your configuration here, go to NetApp Hardware Universe for comprehensive configuration and slot priority information to cable your storage system.

Ethernet host cabling

FAS50 with one 4-port 10/25 GbE I/O module

Steps

1. On each controller, cable ports e2a, e2b, e2c and e2d to the Ethernet host network switches.

10/25 GbE cables

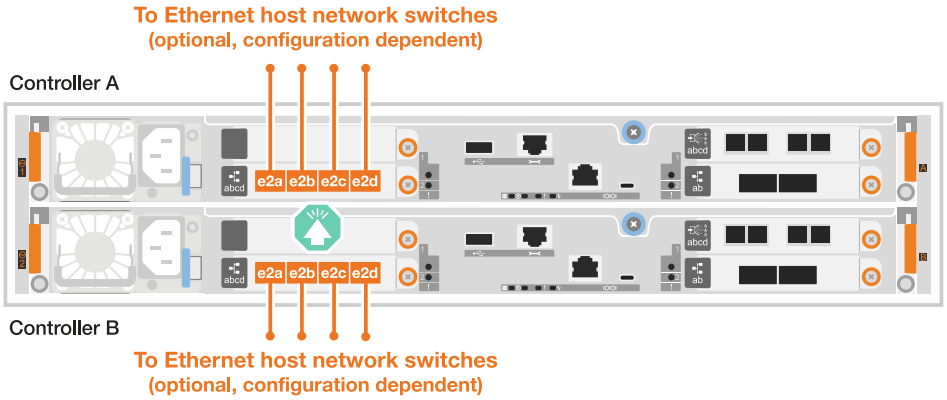

FC host cabling

FAS50 with one 4-port 64 Gb/s FC I/O module

Steps

1. On each controller, cable ports 1a, 1b, 1c and 1d to the FC host network switches.

64 Gb/s FC cables

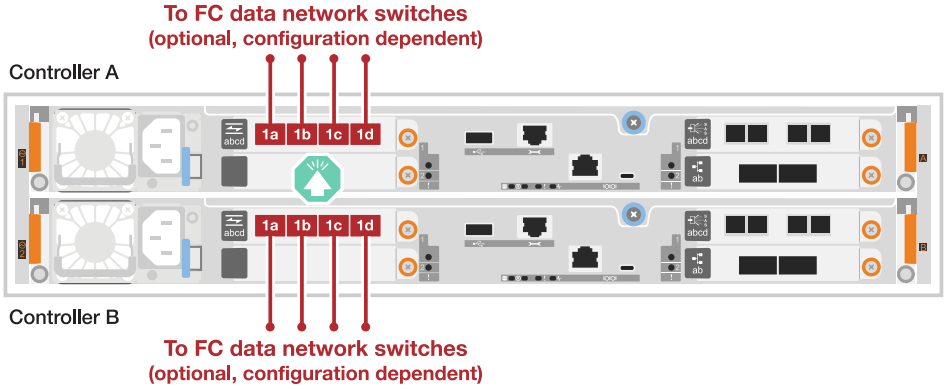

Step 3: Cable the management network connections

Cable the controllers to your management network.

1. Cable the management (wrench) ports on each controller to the management network switches.

1000BASE-T RJ-45 cables

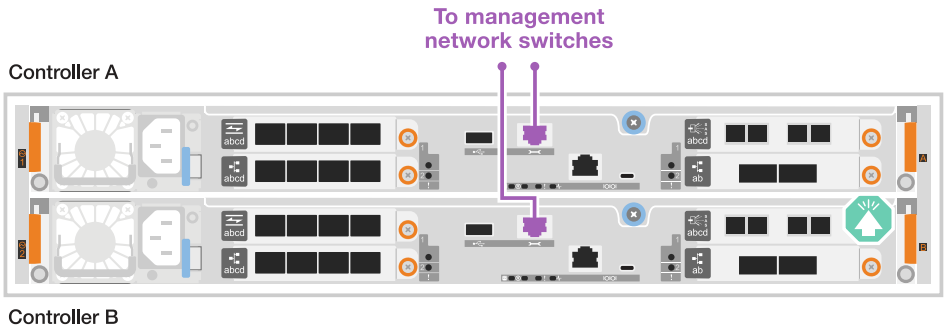

Do not plug in the power cords yet.

Step 4: Cable the shelf connections

The following procedures show how to cable the controllers to one or two DS460C shelves.

About this task

- The cabling examples show DS460C shelves; however, additional SAS shelves are supported, see NetApp Hardware Universe.

  Cabling for other supported SAS shelves is similar. See Install and cable shelves for a new system installation. You can also refer to SAS cabling rues, worksheets, and examples overview.

- For the maximum number of shelves supported for your storage system and all of your cabling options, such as optical and switch-attached, see NetApp Hardware Universe.
- The graphics show controller A cabling in blue and controller B cabling in yellow.
- You use the storage cables that came with your storage system, which could be the following cable type:

  mini-SAS HD cable

Option 1: One DS460C shelf

Cable each controller to each IOM12 module on the DS460C shelf.

Steps

1. Cable controller A to the shelf:
   a. Cable controller A port 3a to IOMA port 1.
   b. Cable controller A port 3d to IOMB port 3.

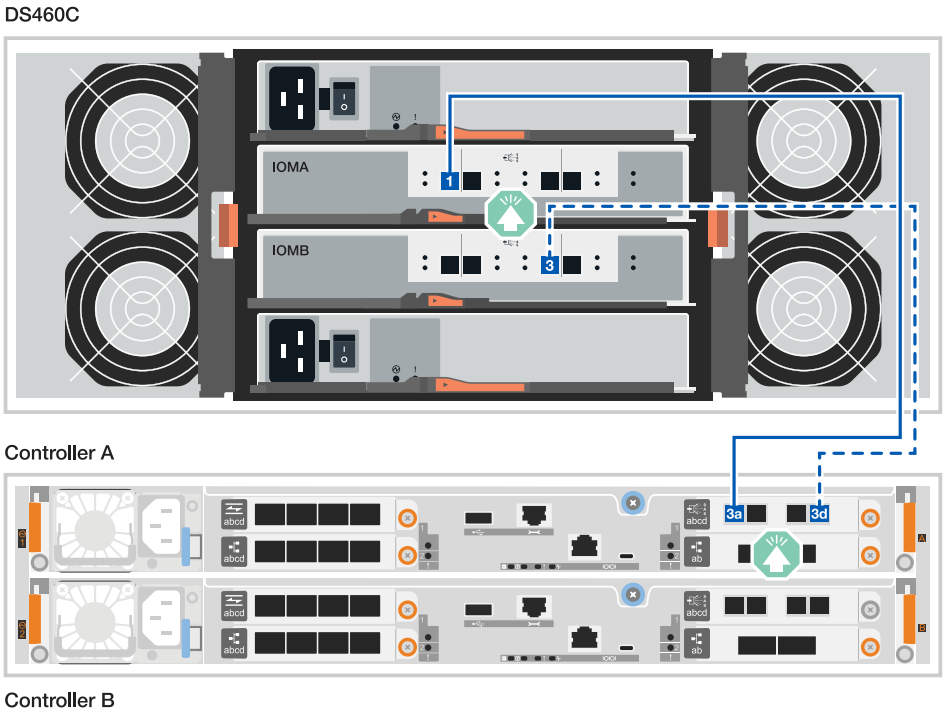

2. Cable controller B to the shelf:
   a. Cable controller B port 3a to IOMB port 1.
   b. Cable controller B port 3d to IOMA port 3.

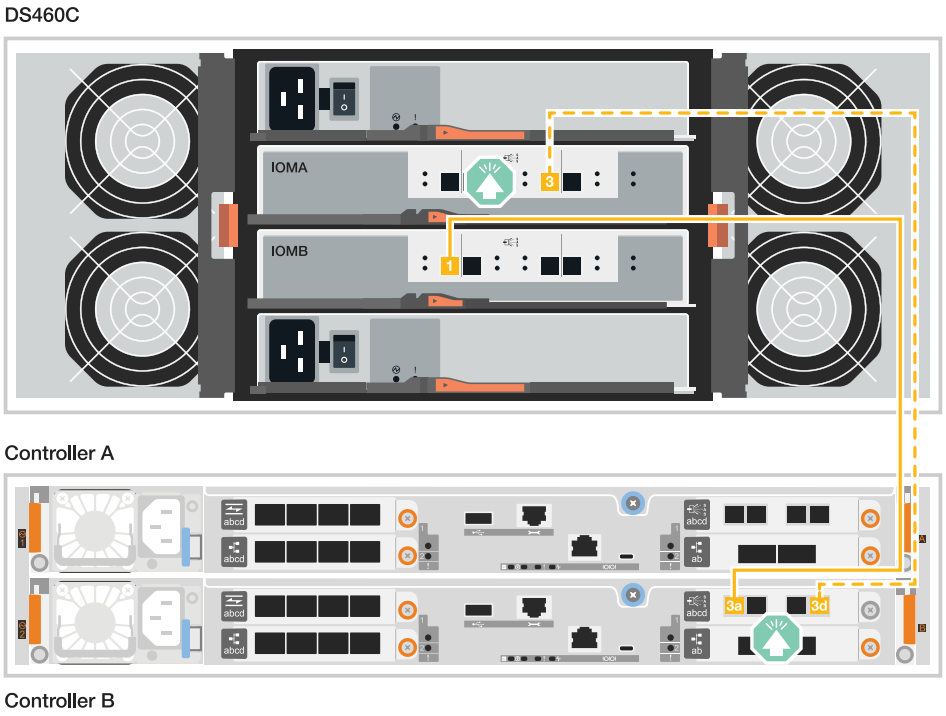

Option 2: Two DS460C shelves

Connect each controller to the IOM12 modules on both DS460C shelves.

1. Cable the shelf-to-shelf connections:
   a. Cable Shelf 1 IOMA port 3 to Shelf 2 IOMA port 1.
   b. Cable Shelf 1 IOMB port 3 to Shelf 2 IOMB port 1.

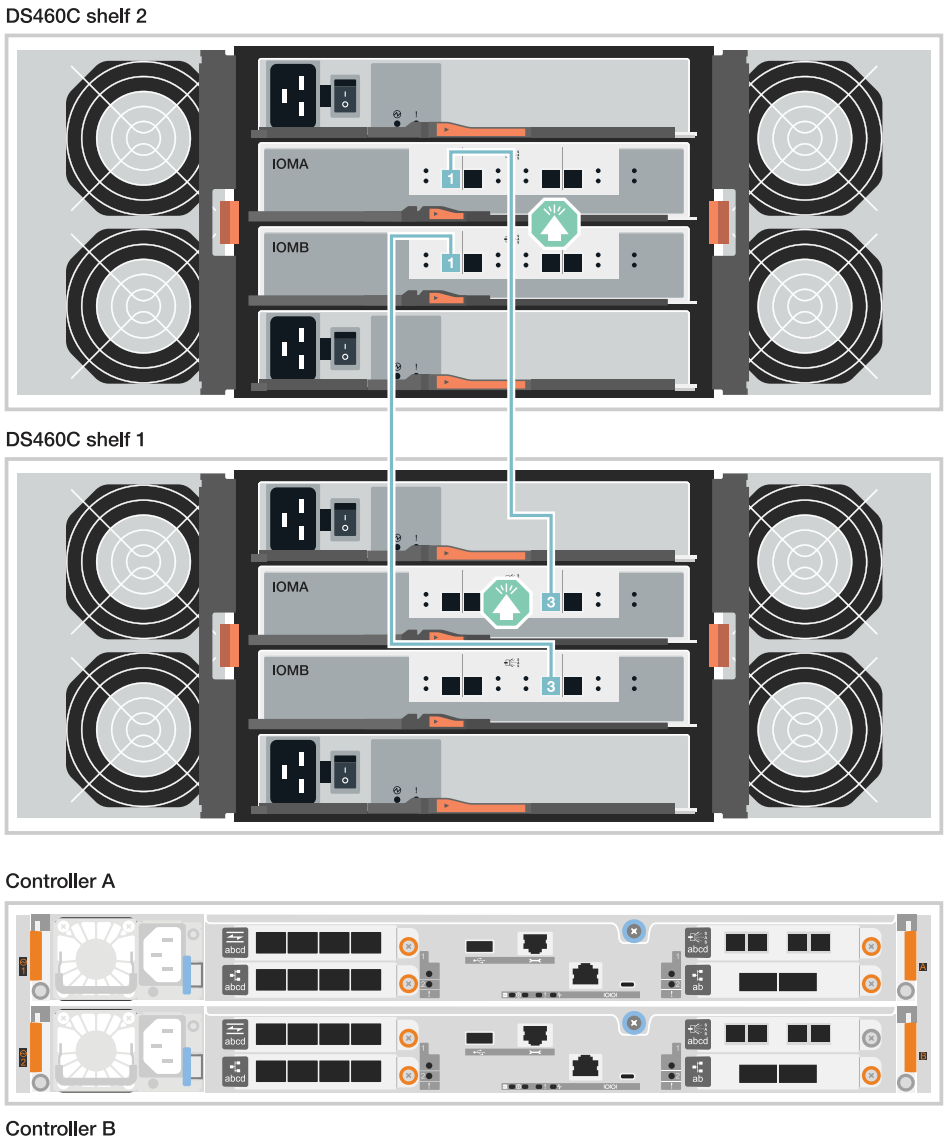

2. Cable controller A to the shelves:
   a. Cable controller A port 3a to shelf 1 IOMA port 1.
   b. Cable controller A port 3d to shelf 2 IOMB port 3.

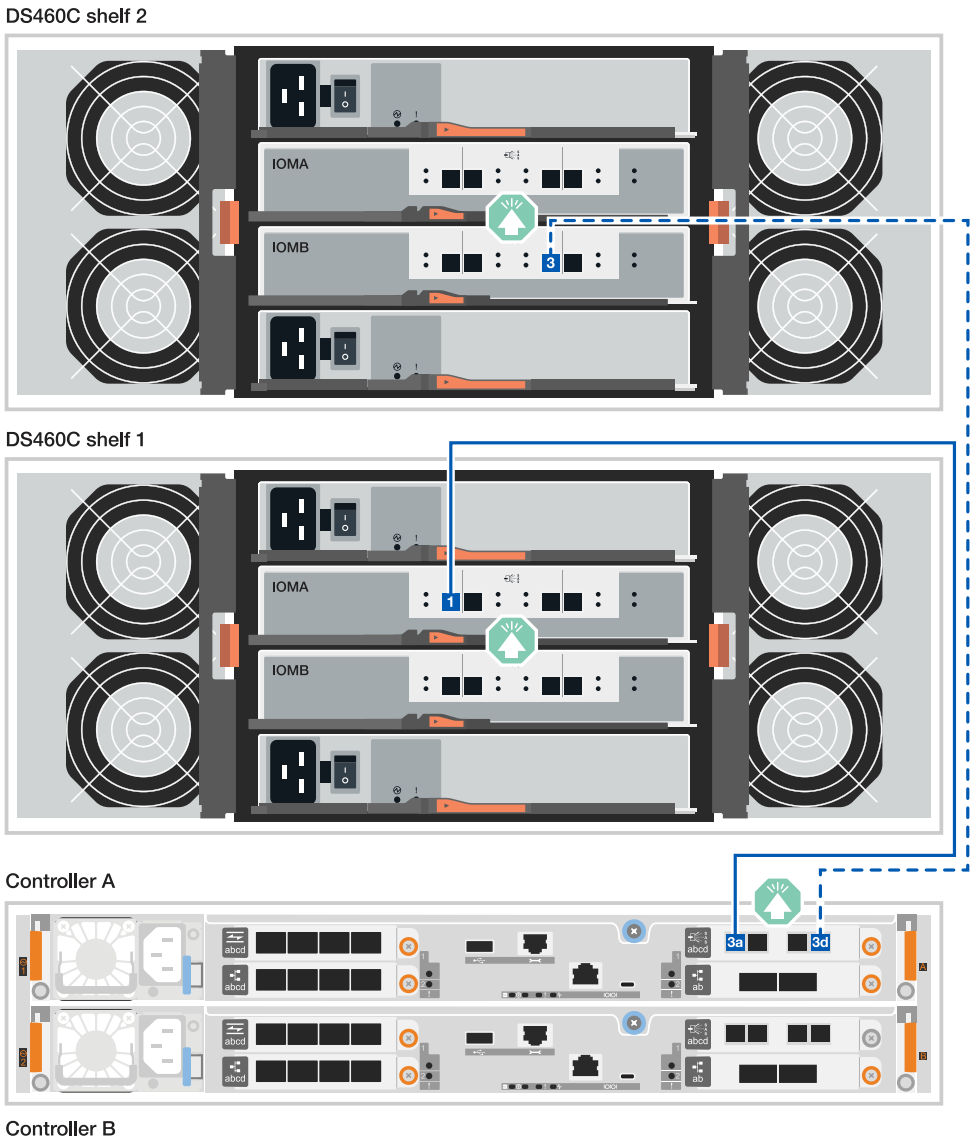

3. Cable controller B to the shelves:
   a. Cable controller B port 3a to shelf 1 IOMB port 1.
   b. Cable controller B port 3d to shelf 2 IOMA port 3.

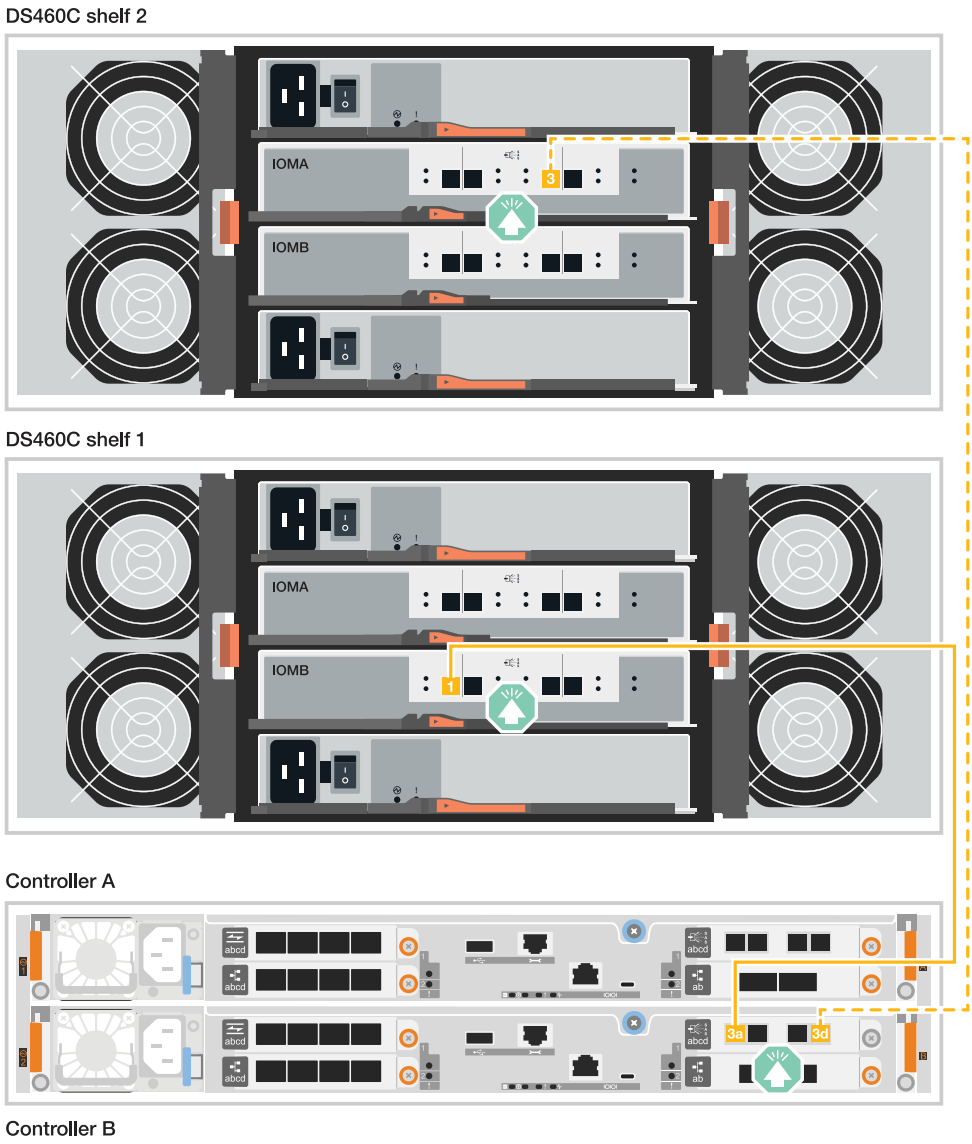

What's next?

After you've cabled the hardware for your storage system, you power on the storage system.

Power on the storage system - FAS50

After you cable the controllers to the network and shelves in your FAS50 storage system, you power on your shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

Each shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed

shelf ID of 00.

- You must power cycle a shelf (turn off the power switch on each of the power supplies of the SAS shelf, wait the appropriate amount of time, and then switch the power back on) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, connecting the power cords to power sources on different circuits, and then turning on the power switch on each of the power supplies (at the rear of the shelf).

   The shelf powers on and boots automatically when powered on.

2. Remove the left end cap to access the orange shelf ID button on the faceplate.

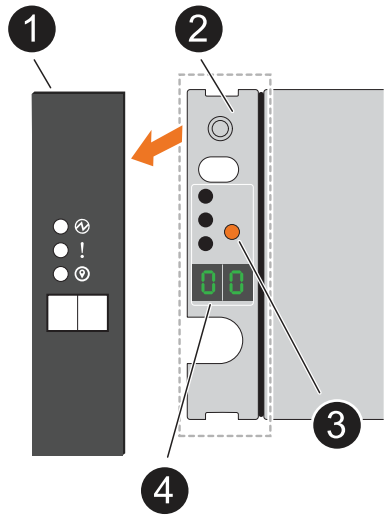

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID button
4	Shelf ID number

3. Change the first number of the shelf ID:
   a. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

      It can take up to 15 seconds for the number to blink. This activates the shelf ID programming mode.

      If the ID takes longer than 15 seconds to blink, press and hold the shelf ID button again, making sure to press it in all the way.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

Each press and release duration can be as short as one second.

The first number continues to blink.

4. Change the second number of the shelf ID:
   a. Press and hold the button until the second number on the digital display blinks.

      It can take up to three seconds for the number to blink.

      The first number on the digital display stops blinking.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

      The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the shelf ID button until the second number stops blinking.

   It can take up to three seconds for the number to stop blinking.

   Both numbers on the digital display start blinking and the amber LED illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.
   a. Turn off the power switch on each of the power supplies.
   b. Wait 10 seconds.
   c. Turn on the power switch on each of the power supplies to complete the power cycle.

      When a power supply is powered on, the bicolored LED should illuminate green.

7. Replace the left end cap.

Step 2: Power on the controllers

After you've powered on your shelves and assigned them unique IDs, power on the storage controllers.

Steps

1. Connect your laptop to the serial console port. This will allow you to monitor the boot sequence when the controllers are powered on.
   a. Set the serial console port on the laptop to 115,200 baud with N-8-1.

      See your laptop's online help for instructions on how to configure the serial console port.

   b. Using the console cable provided with your storage system, connect one end of the console cable to your laptop and the other end to the serial console port on controller A.
   c. Connect the laptop to the switch on the management subnet.

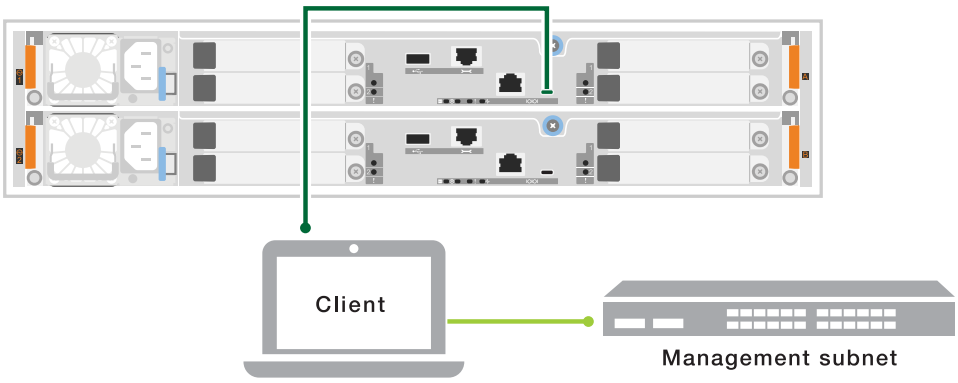

2. Assign a TCP/IP address to the laptop, using one that is on the management subnet.
3. Plug the two power cords into the controller power supplies, and then connect them to power sources on different circuits.

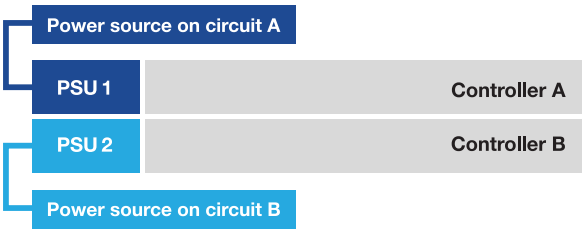

   - The system begins to boot. Initial booting might take up to eight minutes.
   - The LEDs flash on and the fans start, which indicates that the controllers are powering on.
   - The fans might be very noisy when they first start up. The fan noise during start-up is normal.
   - The shelf ID display on the front of the system chassis does not illuminate.
4. Secure the power cords using the securing device on each power supply.

What's next?

After you've powered on your storage system, you set up your cluster.

Maintain

Overview of hardware maintenance - FAS50

Maintain the hardware of your FAS50 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS50 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS50 storage systems, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot the image from the partner node and automatically run the appropriate boot menu option to install the boot image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media- manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot the image from a USB drive and restore the configuration from the partner node.
Caching module	A caching module (Flash Cache module) utilizes high-speed SSDs to store frequently accessed data for quicker retrieval.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Fan	A fan cools the controller and drives.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
NV battery	The non-volatile memory (NV) battery is responsible for providing power to the NVMEM components while data in-flight is being destaged to flash memory after a power loss.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automatic recovery workflow - FAS50

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your FAS50 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1 Review the boot media requirements

Review the requirements for boot media replacement.

2 Shut down the controller

Shut down the controller in your storage system when when you need to replace the boot media.

3 Replace the boot media

Remove the failed boot media from the impaired controller and install the replacement boot media.

4 Restore the image on the boot media

Restore the ONTAP image from the partner controller.

5 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - FAS50

Before replacing the boot media in your FAS50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0M (wrench) port on the impaired controller is working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Review the following requirements.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.

- /cfcard/kmip/certs/client.key file.
- /cfcard/kmip/certs/CA.pem file.

- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - FAS50

Shut down the impaired controller in your FAS50 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - FAS50

The boot media in your FAS50 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

About this task

If needed, you can turn on the platform chassis location (blue) LEDs to aid in physically locating the affected platform. Log into the BMC using SSH and enter the `system location-led on` command.

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

> If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.

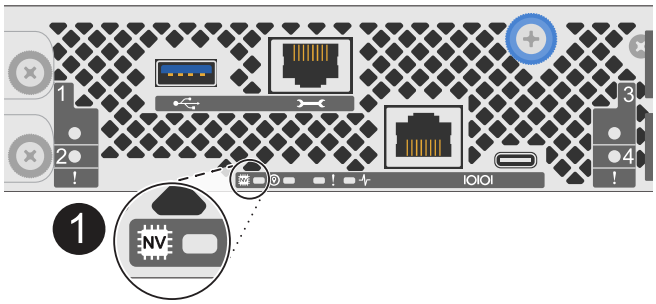

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

> Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

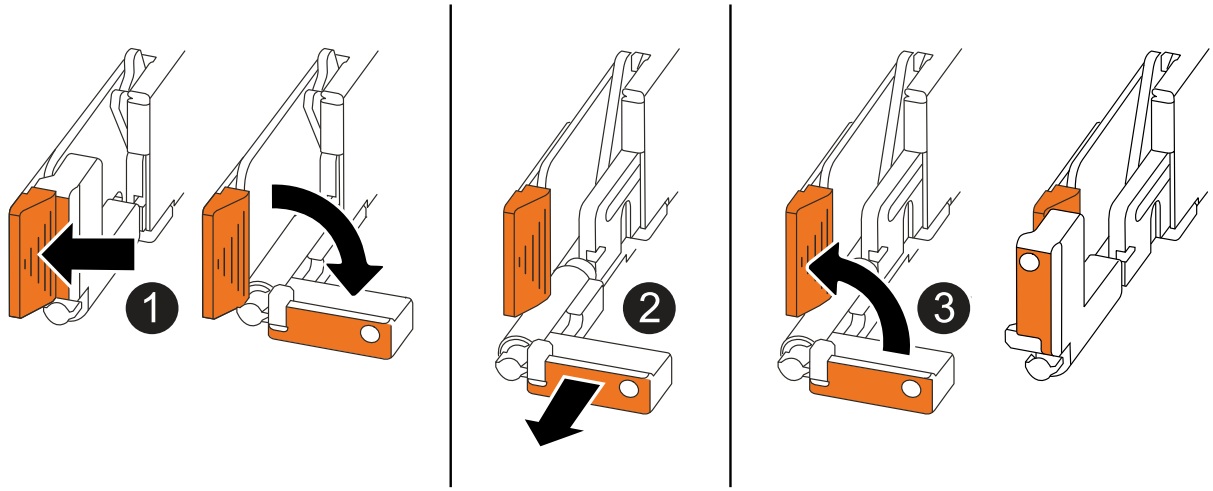

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.
7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:

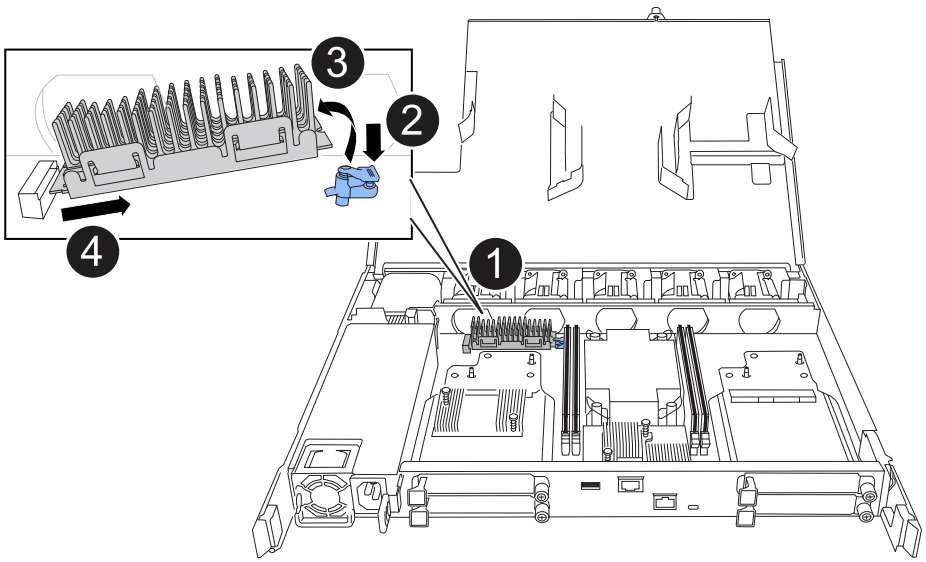

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:
   a. Remove the boot media from its package.
   b. Slide the socket end of the boot media into its socket.
   c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

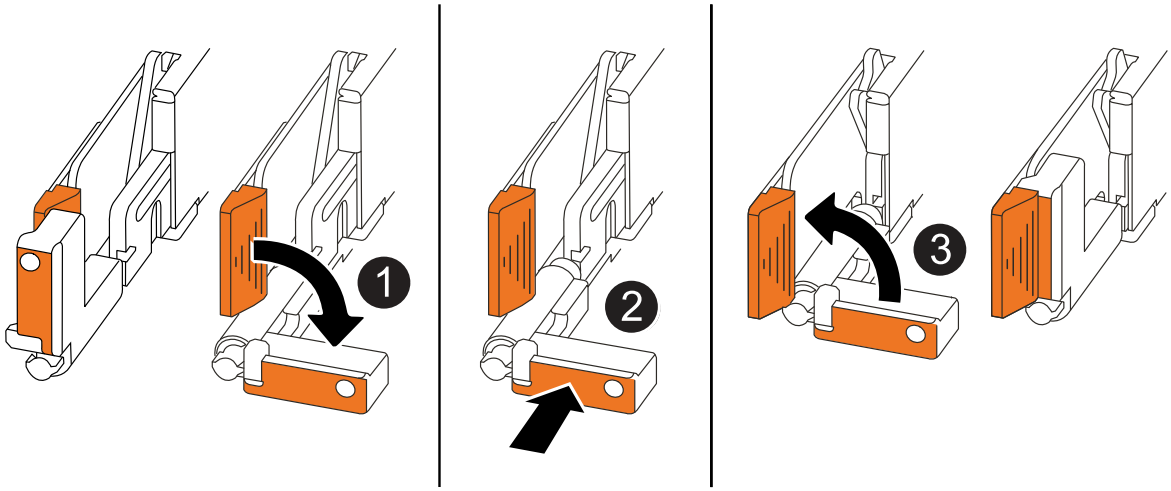

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.

   Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

      The controller boots to the LOADER prompt when fully seated in the chassis. It gets its power from the partner controller.

b. Rotate the controller handles up and lock in place with the tabs.

5. Reconnect the power cord to the PSU on the impaired controller.

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - FAS50

After installing the new boot media device in your FAS50 storage system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media part to NetApp - FAS50

If a component in your FAS50 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - FAS50

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the FAS50 system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

If your system is running in ONTAP 9.17.1 and later, use the automatic boot recovery procedure.

To get started, review the recovery requirements, shut down the controller, replace the boot media, use the USB drive to restore the image, and reapply encryption settings if necessary.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check onboard encryption keys

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the impaired controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the impaired controller and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONTAP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - FAS50

Before replacing the boot media in your FAS50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption support for manual boot media recovery - FAS50

To ensure data security on your FAS50 storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

What's next?

After checking the encryption key support and status on the boot media, you need to shut down the controller.

Shut down the controller for manual boot media recovery - FAS50

Shut down the impaired controller in your FAS50 storage system to prevent data loss and maintain system stability during the manual boot media recovery process.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After shutting down the controller, you need to replace the boot media.

Replace the boot media and prepare for manual boot recovery - FAS50

The boot media in your FAS50 system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

About this task

If needed, you can turn on the platform chassis location (blue) LEDs to aid in physically locating the affected platform. Log into the BMC using SSH and enter the `system location-led on` command.

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.

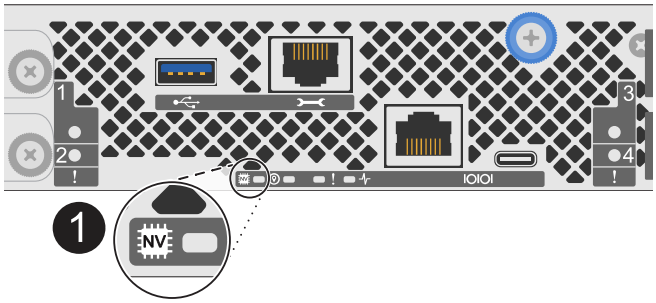

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

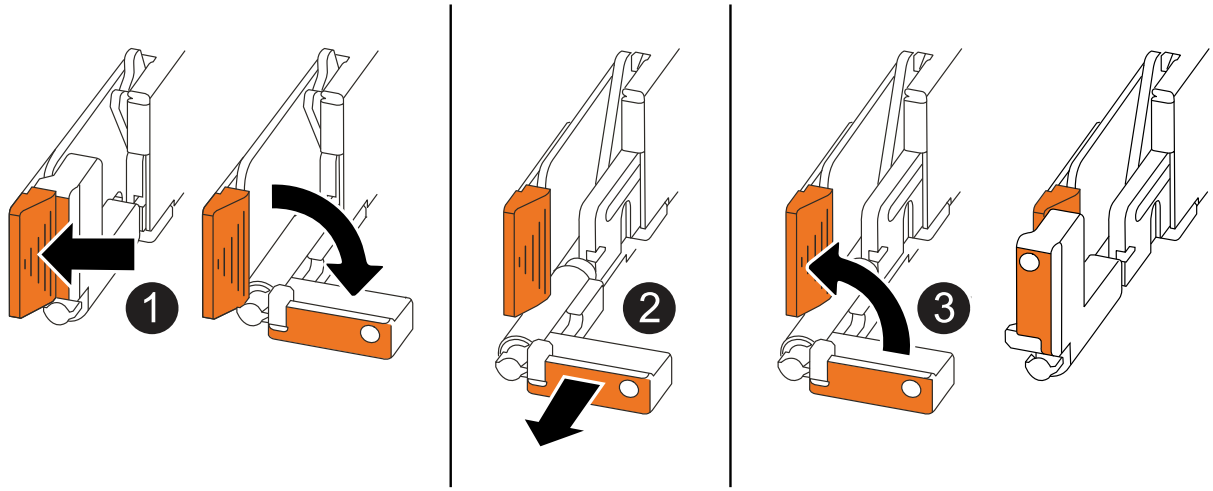

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.
7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:

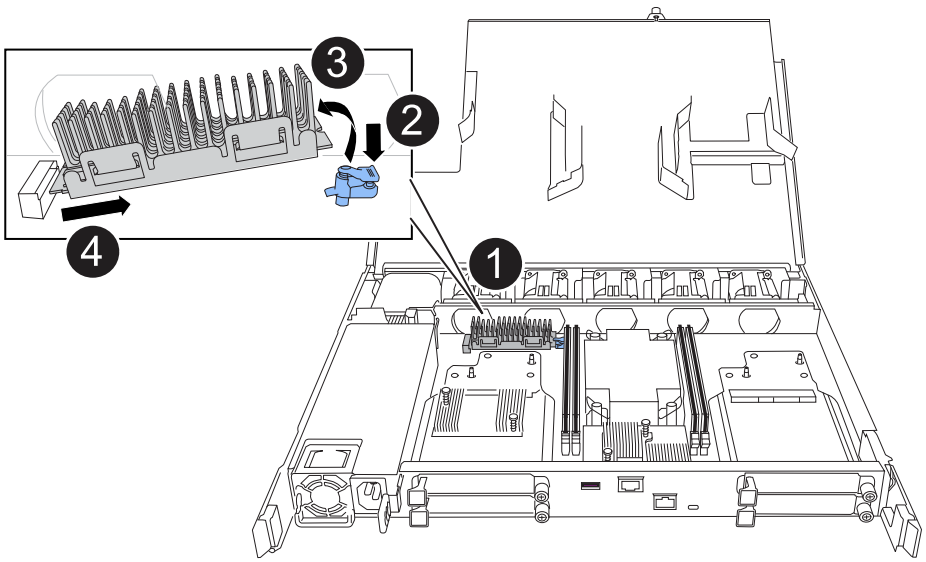

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

3. Install the replacement boot media:
   a. Remove the boot media from its package.
   b. Slide the socket end of the boot media into its socket.
   c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis, but do not reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

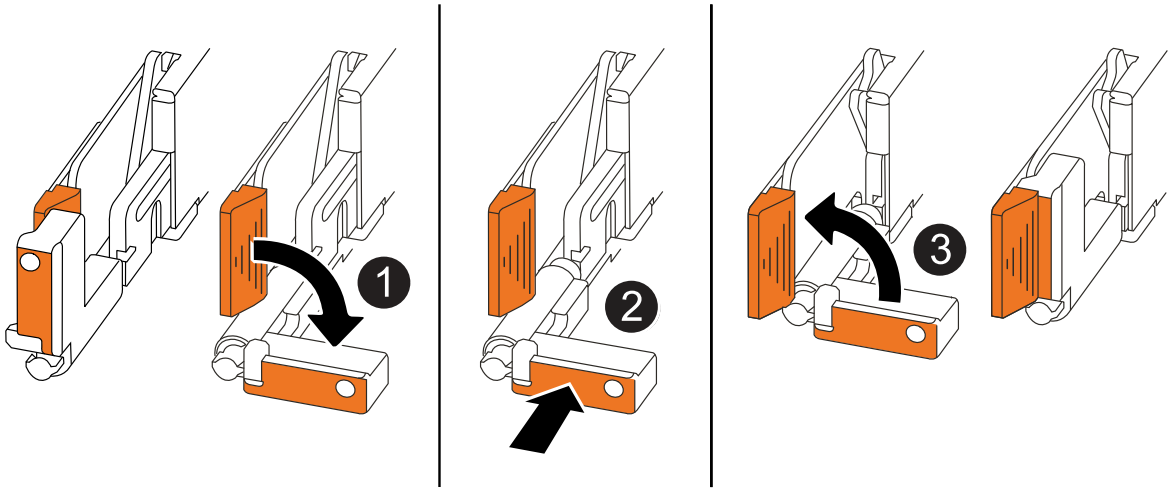

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   > Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.

   > Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

Step 4: Transfer the boot image to the boot media

The replacement boot media that you installed is without an ONTAP image so you need to transfer an ONTAP image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- You must have a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site

- If NVE is supported, download the image with NetApp Volume Encryption, as indicated in the download button.
- If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- You must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image from the Downloads link on the page, to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      The USB flash drive should have the appropriate ONTAP image of what the impaired controller is running.

   c. Remove the USB flash drive from your laptop.
2. Insert the USB flash drive into the USB-A port on the impaired controller.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

3. Fully seat the impaired controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

      The controller boots when fully seated in the chassis. It gets its power from the partner controller.

   b. Rotate the controller handles up and lock in place with the tabs.
4. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

5. Reconnect the power cord to the power supply (PSU) on the impaired controller.

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.

If you are reconnecting a…	Then…
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

After replacing the boot media, you need to boot the recovery image.

Manual boot media recovery from a USB drive - FAS50

After installing the new boot media device in your FAS50 storage system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `...would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

What's next?

After booting the recovery image, you need to restore encryption on the boot media.

Restore encryption keys after manual boot recovery - FAS50

Restore encryption on the replacement boot media in your FAS50 storage system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
**********************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
**********************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

What's next?

After restoring encryption on the boot media, you need to return the failed part to NetApp.

Return the failed part to NetApp - FAS50

If a component in your FAS50 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Hot-swap a caching module - FAS50

You can hot-swap an NVMe SSD caching module (Flash Cache module) of the same capacity from the same or different supported vendor for your FAS50 storage system.

Before you begin

Your storage system must meet certain criteria depending on your situation:

- Your storage system must have the appropriate operating system for the caching module you are installing.
- The replacement caching module must have the same capacity as the failed caching module, but can be from a different supported vendor.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing.

Steps

1. Properly ground yourself.
2. Remove the bezel from the front of the storage system.
3. Locate the failed caching module by the lit amber Attention LED on the front of the caching module.

   A caching module can be in drive bay 0 or 23.

4. Remove the caching module:

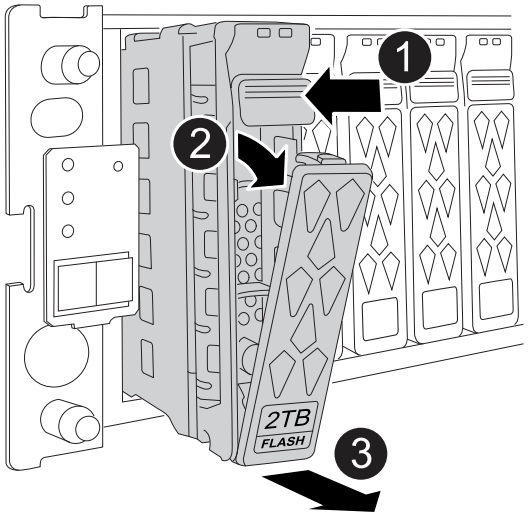

1	Press the release button on the module face to open the cam handle.
2	Rotate the cam handle downward to disengage the module from the midplane.
3	Slide the module out of the drive bay using the cam handle and supporting the module with your other hand. When removing a module, always use two hands to support its weight.

5. Wait a minimum of 70 seconds before inserting the replacement caching module.
6. Install the replacement caching module:
   a. With the cam handle in the open position, use both hands to insert the module.
   b. Gently push until the module stops.
   c. Close the cam handle so that the module is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the module.

7. Verify that the module's activity (green) LED is illuminated.
8. Reinstall the bezel on the front of the storage system.
9. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - FAS50

Follow these workflow steps to replace your FAS50 storage system chassis.

Review the chassis replace requirements

To replace the chassis, you must meet certain requirements.

Shut down the controllers

Shut down the controllers so you can perform maintenance on the chassis.

Replace the chassis

Replacing the chassis includes moving the drive blanks, any caching modules, controllers (with the power supplies), and bezel from the impaired chassis to the new chassis, and swapping out the impaired chassis with the new chassis of the same model as the impaired chassis.

Complete chassis replacement

Verify the HA state of the chassis and return the failed part to NetApp.

Chassis replacement requirements - FAS50

Before replacing the chassis of your FAS50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement chassis, and the necessary tools.

Review the following requirements and considerations.

Requirements

- The replacement chassis must be the same model as the impaired chassis. This procedure is for a like-for-like replacement, not for an upgrade.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Considerations

- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service

outage and a partial outage in a multi-node cluster.

- You can use the chassis replacement procedure with all versions of ONTAP supported by your storage system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, drives, any drive blanks, and controllers to the new chassis.

What's next?

After you've reviewed the requirements to replace the chassis, you need to shut down the controllers.

Shut down the controllers - FAS50

Shut down the controllers in your FAS50 storage system to prevent data loss and ensure system stability when replacing the chassis.

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

What's next?

After you've shut down the controllers, you need to replace the chassis.

Replace the chassis - FAS50

Replace the chassis of your FAS50 storage system when a hardware failure requires it. The replacement process involves removing the controllers, removing any caching modules and drive blanks, installing the replacement chassis, and reinstalling the chassis components.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

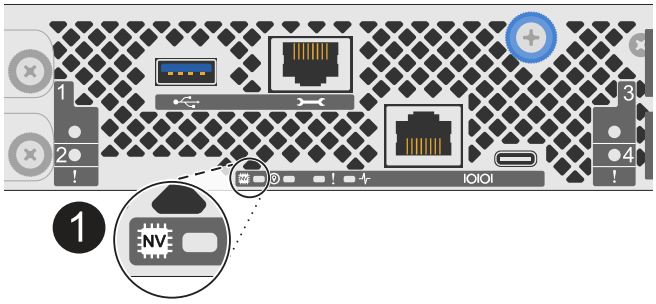

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

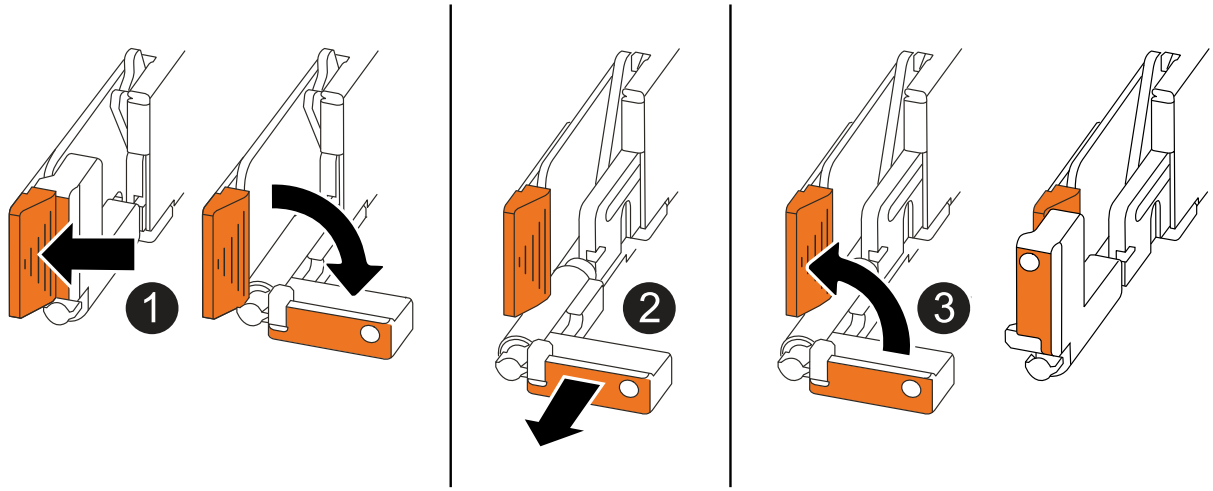

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Repeat these steps for the other controller in the chassis.

Step 2: Remove the caching modules from the impaired chassis

You need to remove any caching modules and drive blanks from the impaired chassis so that later in the procedure you can install them in the replacement chassis.

1. Gently remove the bezel from the front of the storage system.
2. Remove the caching modules and drive blanks:

Keep track of what drive bay each caching module was removed from because they must be installed in the same drive bays in the replacement chassis.

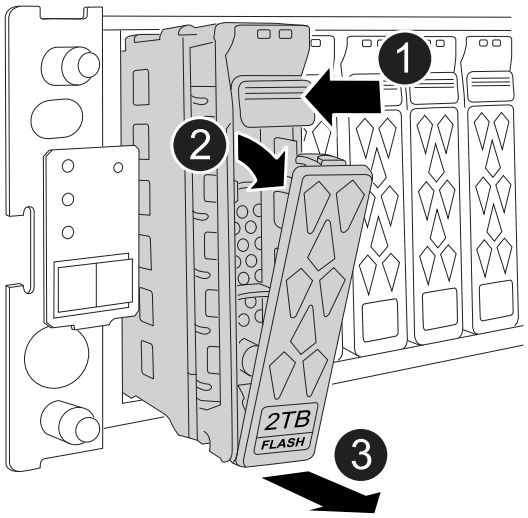

1	Press the release button on the caching module face to open the cam handle.
2	Rotate the cam handle downward to disengage the caching module from the midplane.
3	Slide the caching module out of the drive bay using the cam handle and supporting the caching module with your other hand. When removing a caching module, always use two hands to support its weight.

3. Set the caching modules aside on a static-free cart or table.

Step 3: Replace the chassis from within the equipment rack or system cabinet

You remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis, install the controllers, install any caching modules and drive blanks and then install the bezel.

1. Remove the screws from the impaired chassis mount points.

   Set the screws aside to use later in this procedure.

   If the storage system shipped in a NetApp system cabinet, you must remove additional screws at the rear of the chassis before the chassis can be removed.

2. Using two people or a power lift, remove the impaired chassis from the equipment rack or system cabinet by sliding it off the rails, and then set it aside.
3. Using two people, install the replacement chassis into the equipment rack or system cabinet by sliding it onto the rails.
4. Secure the front of the replacement chassis to the equipment rack or system cabinet using the screws you removed from the impaired chassis.

Step 4: Install the controllers

Install the controllers into the replacement chassis and reboot them.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when installing a controller, and can be used as a reference for the rest of the controller installation steps.

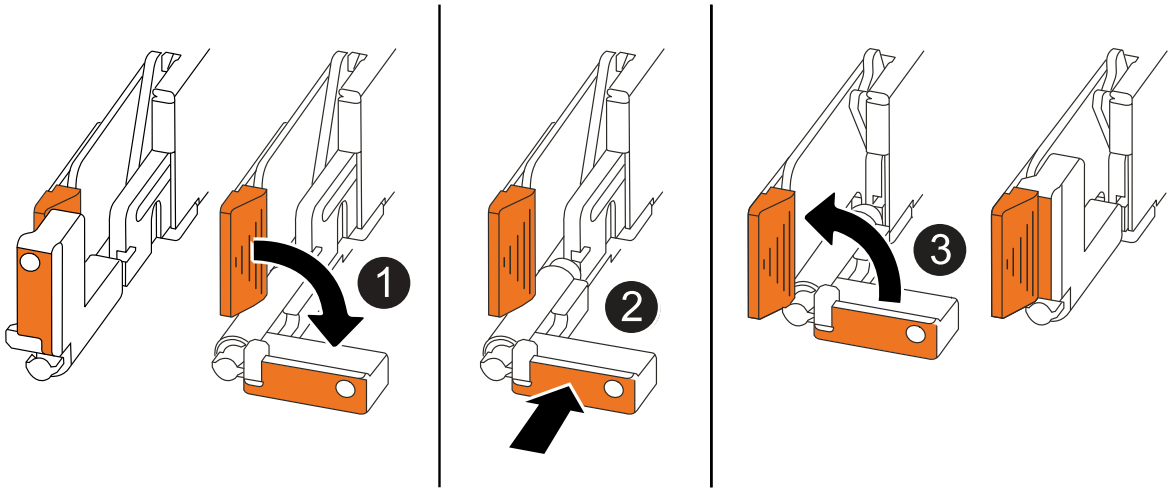

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis and push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

1. Insert one of the controllers into the chassis:
   a. Align the back of the controller with the opening in the chassis.
   b. Firmly push on the handles until the controller meets the midplane and is fully seated in the chassis.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   c. Rotate the controller handles up and lock in place with the tabs.
2. Recable the controller, as needed, except for the power cords.
3. Repeat these steps to install the second controller into the chassis.
4. Install the caching modules and drive blanks you removed from the impaired chassis into the replacement chassis:

   The caching modules and drive blanks must be installed in the same drive bays in the replacement chassis.

   a. With the cam handle in the open position, use both hands to insert the caching module.
   b. Gently push until the caching module stops.
   c. Close the cam handle so that the caching module is fully seated into the midplane and the handle clicks into place.

Be sure to close the cam handle slowly so that it aligns correctly with the face of the caching module.

d. Repeat the process for the remaining caching module, if needed.

5. Install the bezel.
6. Reconnect the power cords to the power supplies (PSU) in the controllers.

   Once power is restored to a PSU, the status LED should be green.

   The controllers begin to boot as soon as the power is restored.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. If controllers boot to the LOADER prompt, reboot the controllers:

   ```
   boot_ontap
   ```

8. Turn AutoSupport back on:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

What's next?

After you've replaced the impaired FAS50 chassis and reinstalled the components into it, you need to complete the chassis replacement.

Complete chassis replacement - FAS50

Verify the HA state of the chassis and then return the failed part to NetApp to complete the final step in the FAS50 chassis replacement procedure.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your storage system configuration.

1. In Maintenance mode, from either controller, display the HA state of the local controller and chassis:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your storage system configuration:
   a. Set the HA state for the chassis:

```
ha-config modify chassis HA-state
```

The value for HA-state should be *ha*.
The value for HA-state can be one of the following:
* ***ha***
* *mcc* (not supported in ASA)

b. Confirm that the setting has changed:

```
ha-config show
```

3. If you have not already done so, recable the rest of your storage system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - FAS50

Follow these workflow steps to replace your controller in your FAS50 storage system.

Review the controller replacement requirements

To replace the controller, you must meet certain requirements.

Shut down the impaired controller

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

Replace the controller

Replacing the controller includes removing the impaired controller, moving FRU components to the replacement controller, installing the replacement controller in the chassis, setting the time and date, and then recabling.

Restore and verify the system configuration

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Give back the controller

Transfer the ownership of storage resources back to the replacement controller.

6 Complete controller replacement

Verify the LIFs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - FAS50

Before replacing the controller in your FAS50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements and considerations for the controller replacement procedure.

Requirements

- All shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.
- You must replace a controller with a controller of the same model type. You cannot upgrade your system by just replacing the controller.
- You cannot change any drives or shelves as part of this procedure.
- You must always capture the controller's console output to a text log file.

  The console output provides you with a record of the procedure you can use to troubleshoot issues you might encounter during the replacement process.

Considerations

It is important that you apply the commands in this procedure to the correct controller:

- The *impaired* controller is the controller that is being replaced.
- The *replacement* controller is the new controller that is replacing the impaired controller.
- The *healthy* controller is the surviving controller.

What's next?

After you've reviewed the requirements to replace the impaired controller, you need to shut down the impaired controller.

Shut down the impaired controller - FAS50

Shut down the impaired controller in your FAS50 storage system to prevent data loss and ensure system stability when replacing the controller.

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After you've shut down the impaired controller, you need to replace the controller.

Replace the controller - FAS50

Replace the controller in your FAS50 storage system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.

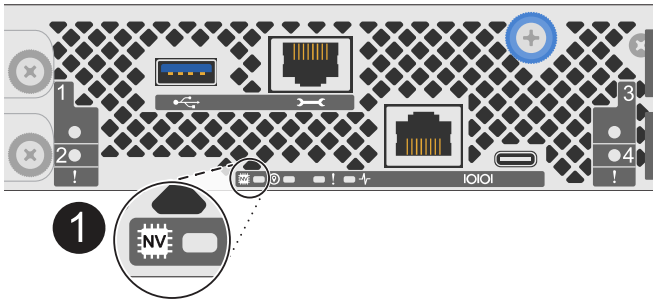

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

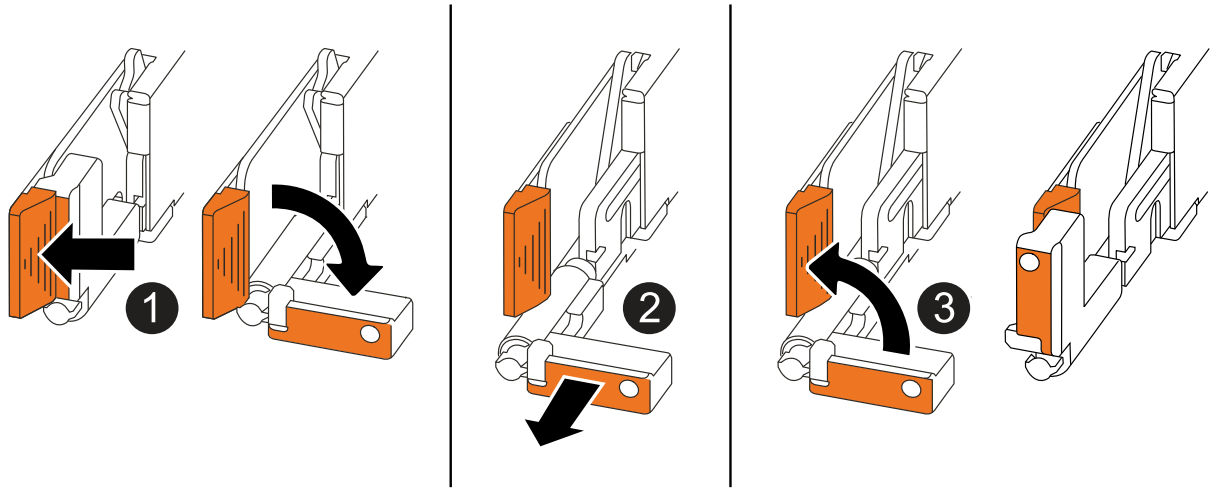

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Move the power supply

Move the power supply (PSU) to the replacement controller.

1. Move the PSU from the impaired controller:

   Make sure the left side controller handle is in the upright position to allow you access to the PSU.

Option 1: Move an AC PSU

To move an AC PSU, complete the following steps.

Steps

1. Remove the AC PSU from the impaired controller:

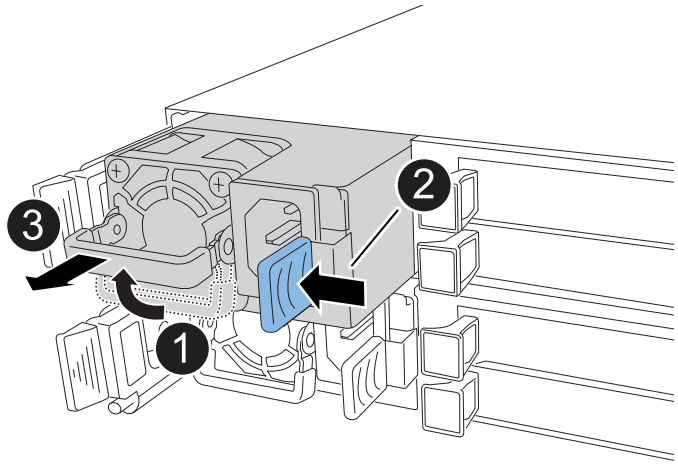

1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

2. Insert the PSU into the replacement controller:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.

Option 2: Move a DC PSU

To move a DC PSU, complete the following steps.

Steps

1. Remove the DC PSU from the impaired controller:

a. Rotate the handle up, to its horizontal position, and then grasp it.
b. With your thumb, press the terracotta tab to release the locking mechanism.
c. Pull the PSU out of the controller while using your other hand to support its weight.

> The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.

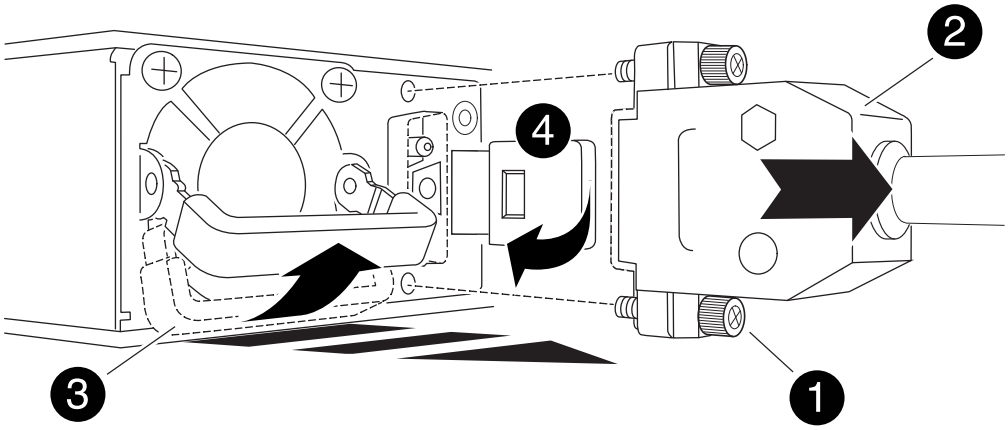

1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

2. Insert the PSU into the replacement controller:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently slide the PSU into the controller until the locking tab clicks into place.

      A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

      > To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.

Step 3: Move the fans

Move the fans to the replacement controller.

1. Remove one of the fans from the impaired controller:

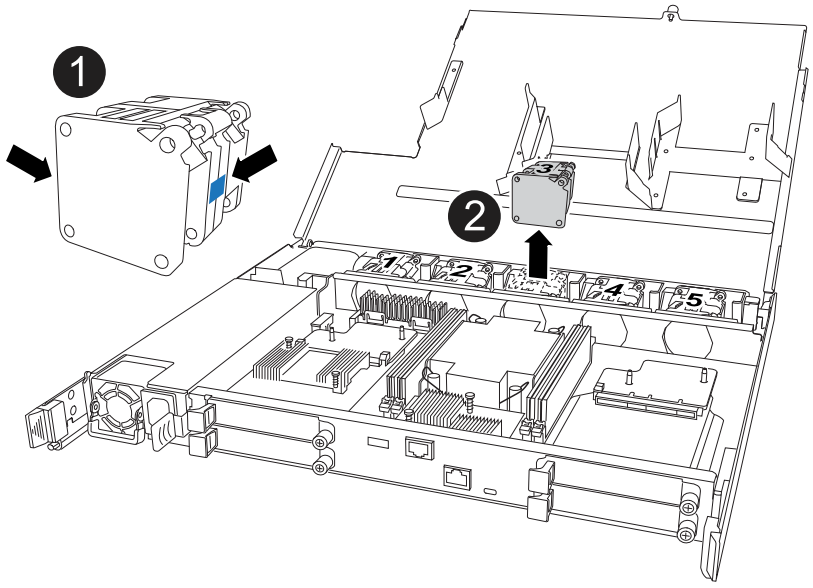

1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

2. Insert the fan into the replacement controller by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.
3. Repeat these steps for the remaining fans.

Step 4: Move the NV battery

Move the NV battery to the replacement controller.

1. Remove the NV battery from the impaired controller:

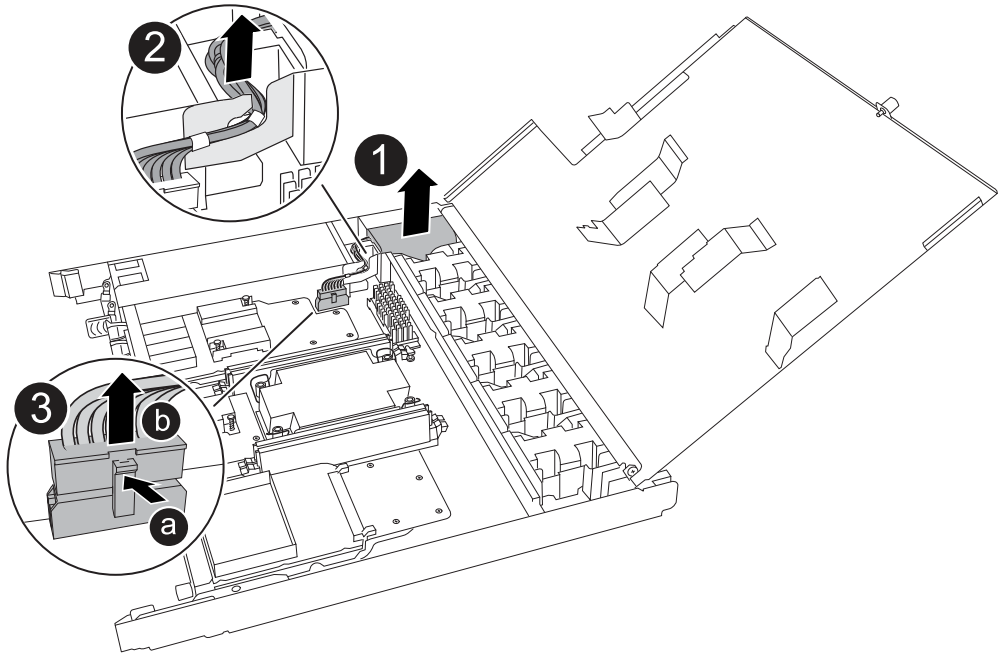

1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

2. Install the NV battery into the replacement controller:
   a. Plug the wiring connector into its socket.
   b. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   c. Place the NV battery into the compartment.

      The NV battery should sit flush in its compartment.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller.

If you have DIMM blanks, you do not need to move them, the replacement controller should come with them installed.

1. Remove one of the DIMMs from the impaired controller:

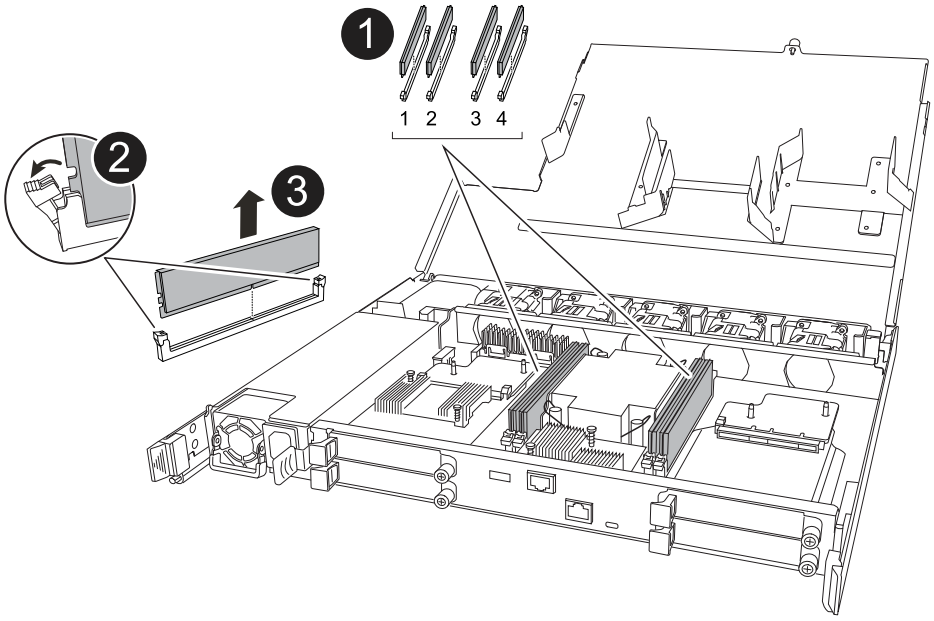

1	DIMM slot numbering and positions. Depending on your storage system model, you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller in the proper orientation. • Eject the DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

2. Install the DIMM in the replacement controller:
   a. Make sure that the DIMM ejector tabs on the connector are in the open position.
   b. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

      When inserted correctly, the DIMM goes in easily but fits tightly in the slot. If not, reinsert the DIMM.

   c. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
   d. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.
3. Repeat these steps for the remaining DIMMs.

Step 6: Move the boot media

Move the boot media to the replacement controller.

1. Remove the boot media from the impaired controller:

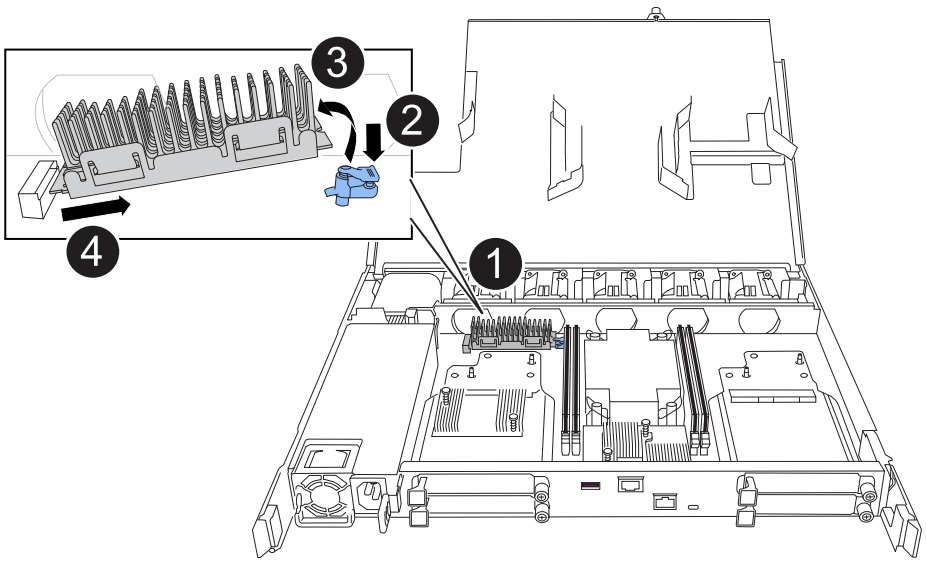

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

2. Install the boot media into the replacement controller:
   a. Slide the socket end of the boot media into its socket.
   b. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 7: Move the I/O modules

Move the I/O modules and any I/O blanking modules to the replacement controller.

1. Unplug cabling from one of the I/O modules.

   Make sure to label the cables so that you know where they came from.

2. Remove the I/O module from the impaired controller:

   Make sure that you keep track of which slot the I/O module was in.

   If you are removing the I/O module in slot 4, make sure the right side controller handle is in the upright position to allow you access to the I/O module.

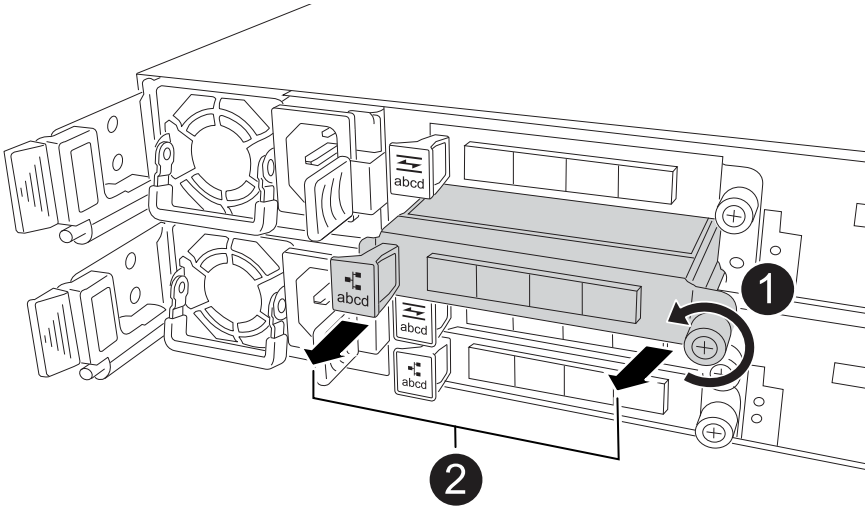

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

3. Install the I/O module into the replacement controller:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
4. Repeat these steps to move the remaining I/O modules and any I/O blanking modules to the replacement controller.

Step 8: Install the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

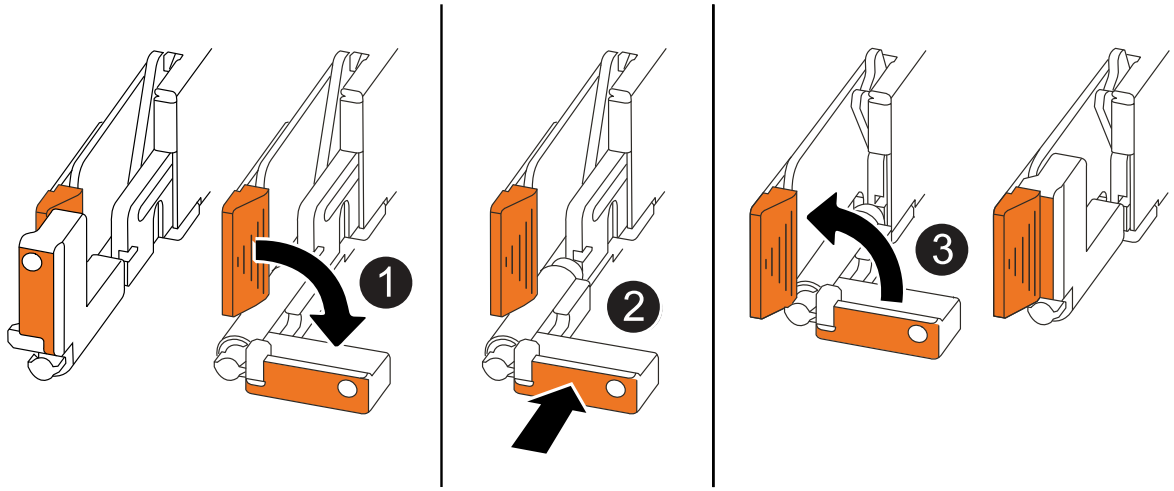

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Take the controller to the LOADER prompt by pressing CTRL-C to abort AUTOBOOT.
6. Set the time and date on the controller:

   Make sure you are at the controller's LOADER prompt.

   a. Display the date and time on the controller:

      ```
      show date
      ```

      Time and date default is in GMT. You have the option to display in local time and in 24hr mode.

   b. Set the current time in GMT:

      ```
      set time hh:mm:ss
      ```

      You can get the current GMT from the healthy node:

      ```
      date -u
      ```

   c. Set the current date in GMT:

      ```
      set date mm/dd/yyyy
      ```

      You can get the current GMT from the healthy node:
      ```
      date -u
      ```
7. Recable the controller as needed.
8. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

After you've replaced the impaired FAS50 controller, you need to restore the system configuration.

Restore and verify the system configuration - FAS50

Verify that the controller's HA configuration is active and functioning correctly in your FAS50 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

You must verify the `HA` state of the controller and, if necessary, update the state to match your storage system configuration.

1. Boot to maintenance mode:

   ```
   boot_ontap maint
   ```

   a. Enter `y` when you see *Continue with boot?*.

      If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

5. If the displayed system state of the controller does not match your storage system configuration, set the `HA` state for the controller:

   ```
   ha-config modify controller ha
   ```

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed:

   ```
   ha-config show
   ```

Step 2: Verify disk list

1. Verify that the adapter lists the paths to all disks:

   ```
   storage show disk -p
   ```

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode:

```
halt
```

What's next?

After you've restored and verified the system configuration for your FAS50 system, you need to give back the controller.

Give back the controller - FAS50

Return control of storage resources to the replacement controller so your FAS50 system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption, Onboard Key Manager (OKM) encryption, or External Key Manager (EKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`

   If you encounter errors, contact NetApp Support.

9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

c. Verify status of the keys: `security key-manager key query -restored false`

   The output should show no results when when properly synchronized.

   If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

External key manager (EKM)

Reset encryption and return the controller to normal operation.

Steps

1. If the root volume is encrypted with External Key Manager and the console cable is connected to the replacement node, enter `boot_ontap menu` and select option `11`.
2. If these questions appear, answer `y` or `n` as appropriate:

   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}

   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}

   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}

   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}

   Do you know the KMIP server address? {y/n}

   Do you know the KMIP port? {y/n}

   Contact NetApp Support if you have issues.

3. Supply the information for:
   - The client certificate (client.crt) file contents
   - The client key (client.key) file contents
   - The KMIP server CA(s) (CA.pem) file contents
   - The IP address for the KMIP server
   - The port for the KMIP server
4. Once the system processes, you see the Boot Menu. Select '1' for normal boot.
5. Check the takeover status: `storage failover show`
6. Ensure any core dumps on the repaired node are saved by going to advanced mode `set -privilege advanced` and then `run local partner nosavecore`.
7. Return the impaired controller to normal operation by giving back its storage: `storage failover`

`giveback -ofnode impaired_node_name`

8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
9. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

After you've transferred the ownership of storage resources to the replacement controller, you need to complete the controller replacement procedure.

Complete controller replacement - FAS50

To complete the controller replacement for your FAS50 system, first restore the NetApp Storage Encryption configuration (if necessary) and install the required licenses on the new controller. Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, register the new controller's serial number and then return the failed part to NetApp.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

  Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The licenses keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.
- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs, register the serial number, and check cluster health

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS50

Replace a DIMM in your FAS50 storage system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- All other components in the storage system must be working correctly; if not, contact NetApp Support before continuing.

- You must replace the failed FRU component with a replacement FRU component you received from your provider.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

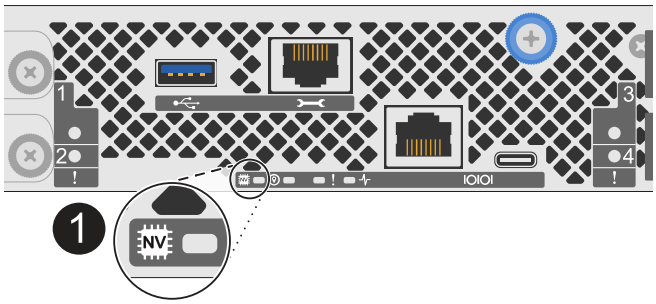

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

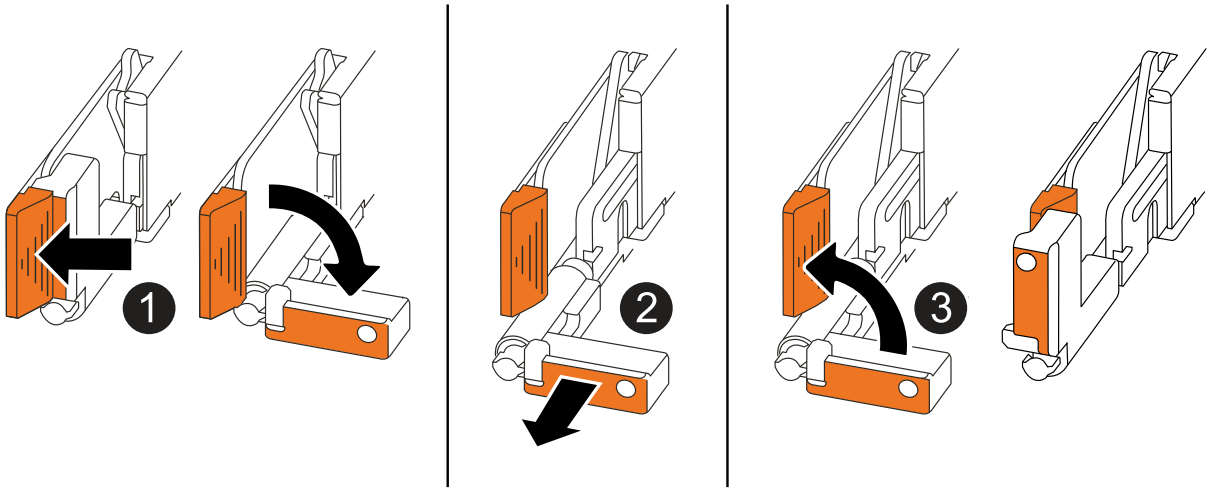

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace a DIMM

To replace a DIMM, locate the faulty DIMM inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller and identify the faulty DIMM.

> Consult either the Netapp Hardware Universe or the FRU map on the cover of the controller for exact DIMM locations.

3. Remove the faulty DIMM:

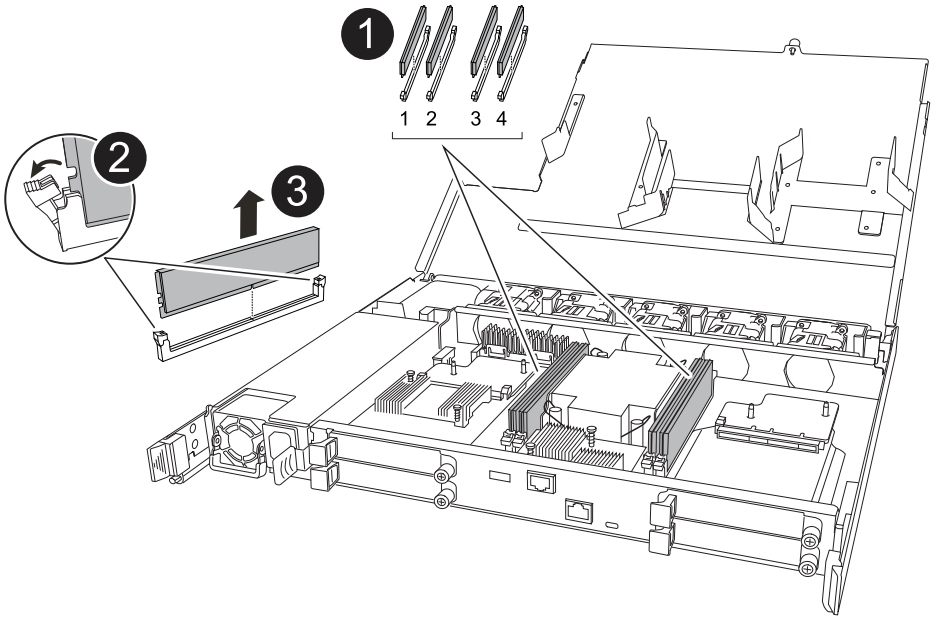

1	DIMM slot numbering and positions. Depending on your storage system model you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM using the same orientation. • Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

4. Install the replacement DIMM:
   a. Remove the replacement DIMM from its antistatic shipping bag.
   b. Make sure that the DIMM ejector tabs on the connector are in the open position.
   c. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

      When inserted correctly, the DIMM goes in easily but fits tightly in the slot. Reinsert the DIMM if you feel it is not inserted correctly.

   d. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
   e. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

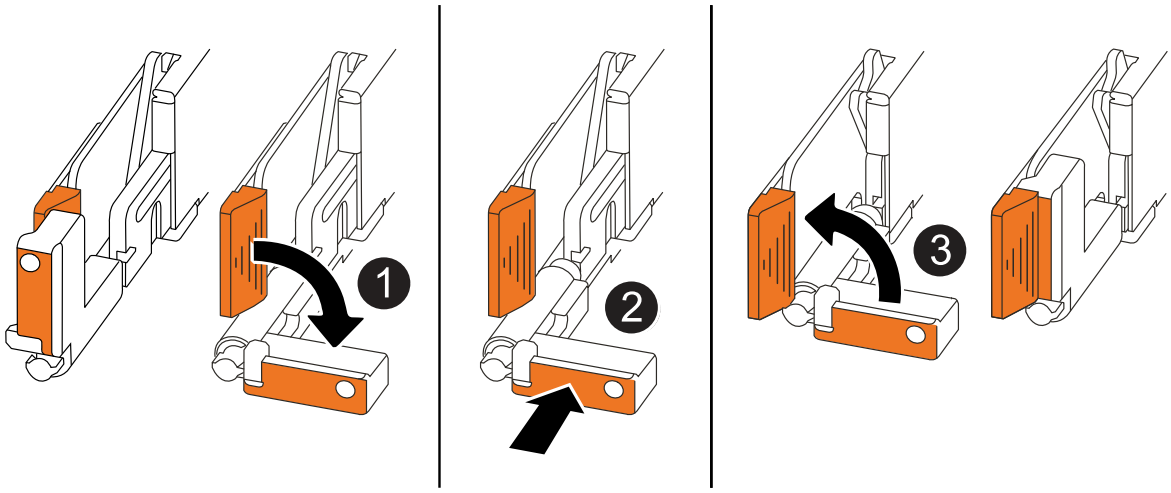

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

> Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

> Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

> Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

> The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a fan module - FAS50

Replace a fan module in your FAS50 system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

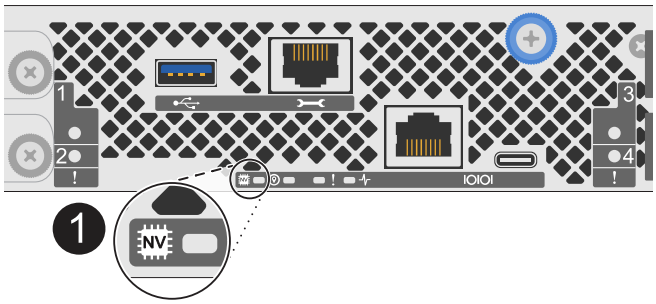

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

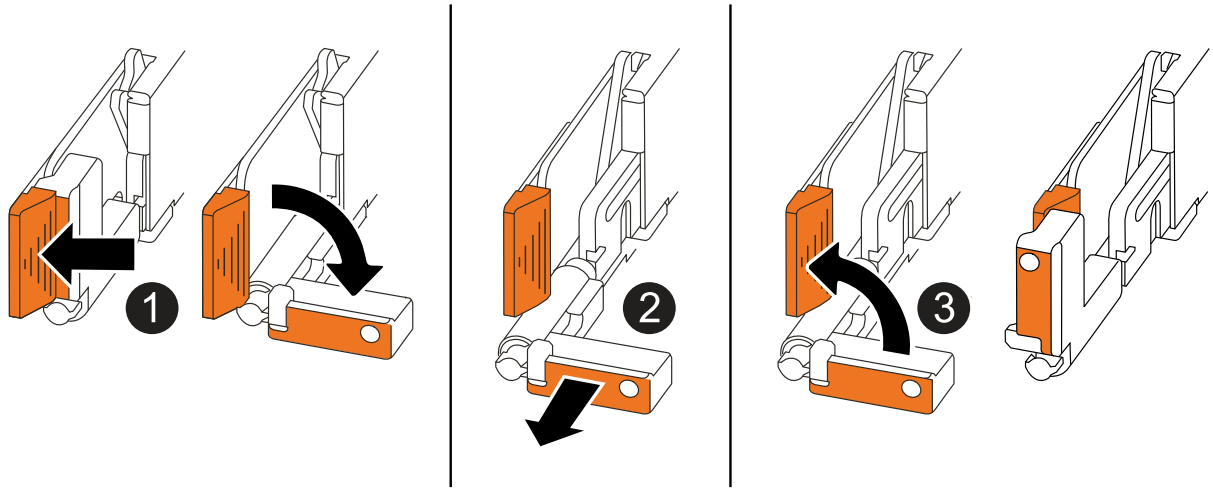

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace fan

To replace a fan, remove the failed fan and replace it with a new fan.

Steps

1. Identify the fan that you must replace by checking the console error messages.
2. Remove the failed fan:

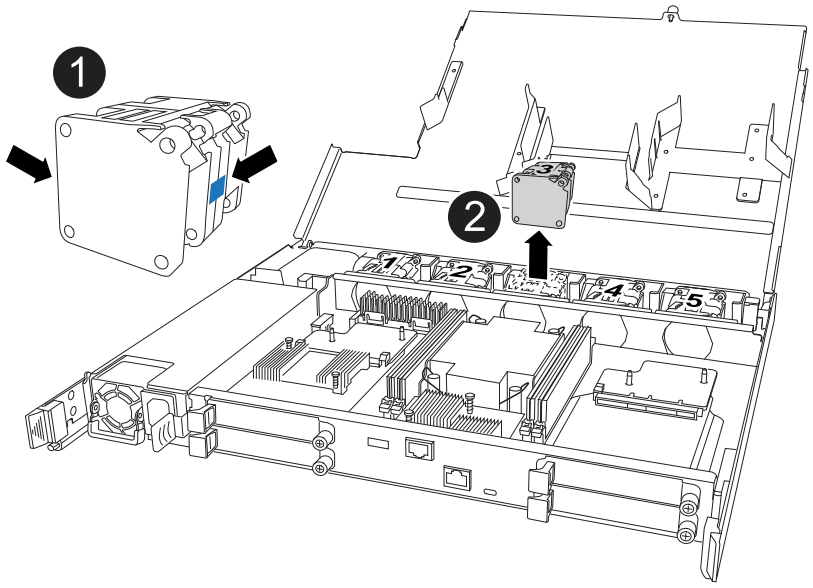

1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

3. Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

Step 4: Reinstall the controller module

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

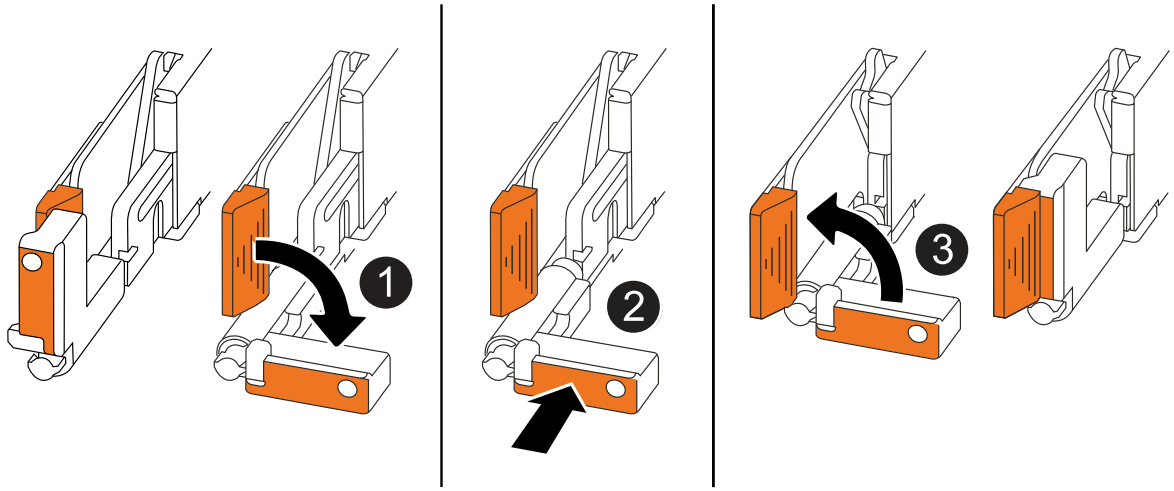

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Overview of I/O module maintenance - FAS50

The FAS50 storage systems offer flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding, hot-swapping, or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your storage system with the same type of I/O module, or with a different type of I/O module. You can hot-swap a cluster and HA I/O module when your storage system meets specific requirements. You can also add an I/O module to a storage system with available slots.

- Add an I/O module

  Adding additional I/O modules can improve redundancy, helping to ensure that the storage system remains operational even if one I/O module fails.

- Hot-swap a cluster and HA I/O module

Hot-swapping a failed cluster and HA I/O module can restore the storage system to its optimal operating state. Hot-swapping is done without having to manually take over the impaired controller.

To use this procedure, your storage system must be running ONTAP 9.17.1 or later and meet specific system requirements.

- Replace an I/O module

  Replacing a failing I/O module can restore the storage system to its optimal operating state.

Add an I/O module - FAS50

Add an I/O module to your FAS50 storage system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your FAS50 storage system when there are slots available or when all slots are fully populated (by removing an existing I/O module and installing a new one in its place).

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller module

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, remove the I/O blanking module from the target slot.

   Unused I/O slots should have blanking module installed to prevent possible thermal issues and for EMC compliance.

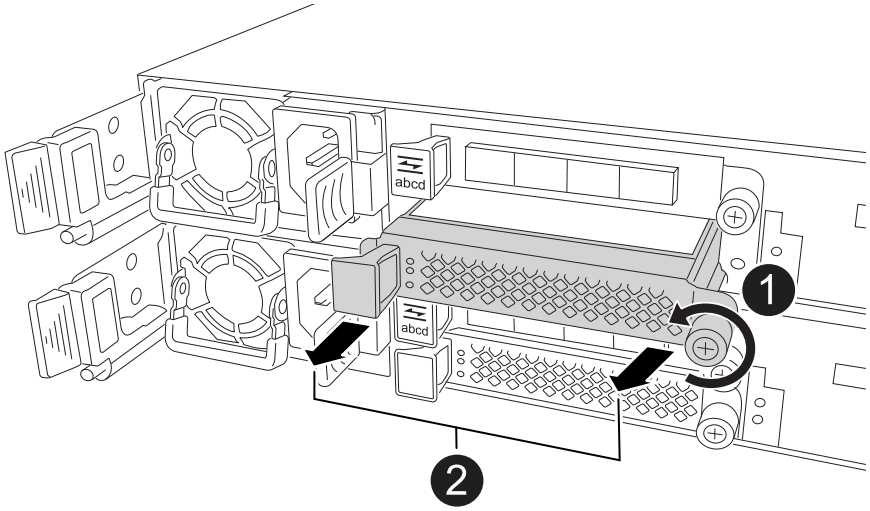

1	On the I/O blanking module, turn the thumbscrew counterclockwise to loosen.
2	Pull the I/O blanking module out of the controller using the tab on the left and the thumbscrew.

3. Install the new I/O module:
   a. Align the I/O module with the edges of the controller slot opening.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O module.

   c. Turn the thumbscrew clockwise to tighten.
4. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.

5. Reboot the impaired controller from the LOADER prompt: `bye`

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

6. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name.
```

7. Repeat these steps to add an I/O module to the other controller.
8. Restore automatic giveback from the console of the healthy controller:

```
storage failover modify -node local -auto-giveback true
```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation: +

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, unplug any cabling on the target I/O module.
3. Remove the target I/O module from the controller:

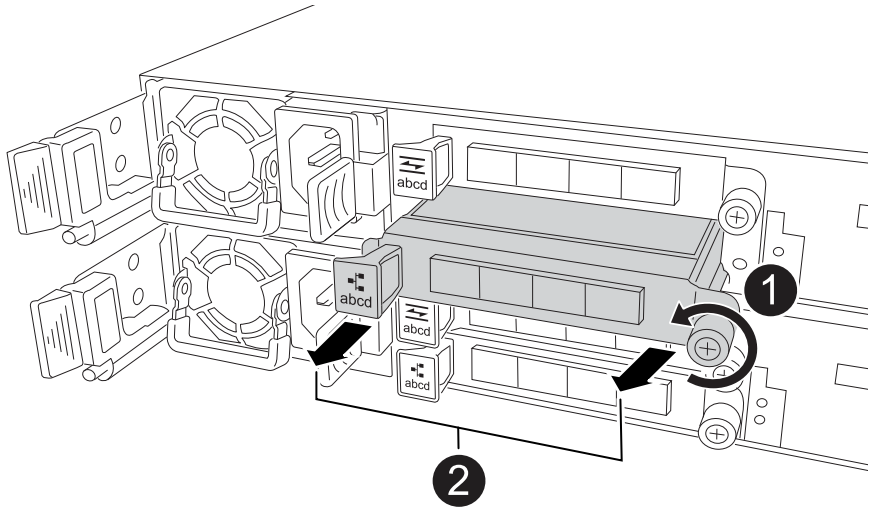

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the new I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
5. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
6. Repeat the I/O module remove and install steps to add any additional I/O modules in the controller.
7. Reboot the impaired controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.
8. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```
9. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```
10. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```
11. If you installed a NIC module, specify the usage mode for each port as *network*:

    ```
    storage port modify -node node_name -port port_name -mode network
    ```
12. Repeat these steps for the other controller.

Hot-swap the I/O module used for cluster and HA traffic - FAS50

The cluster and HA I/O module supports interconnects for clustering and high-availability. You can hot-swap the module in your FAS50 storage system when the module fails and if your storage system meets specific requirements.

To hot-swap a module, you ensure your storage system meets the procedure requirements, prepare the storage system and I/O module in slot 4, hot-swap the failed module for an equivalent one, bring the replacement module online, restore the storage system to normal operation, and return the failed module to NetApp.

About this task

- Hot-swapping the cluster and HA I/O module means that you do not have to perform a manual takeover; the impaired controller (the controller with the failed cluster and HA I/O module) has automatically taken over the healthy controller.

  When the impaired controller has taken over the healthy controller, the only way to recover without an outage is to hot-swap the module.

- It is critical to apply the commands to the correct controller when you are hot-swapping the cluster and HA I/O module:
  - The *impaired controller* is the controller on which you are hot-swapping the cluster and HA I/O module and it is the controller that has taken over the healthy controller.
  - The *healthy controller* is the HA partner of the impaired controller and it is the controller that was taken over by the impaired controller.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, make sure your storage system meets all requirements.

> If your storage system does not meet all requirements, you must use the replace an I/O module procedure.

- Your storage system must be running ONTAP 9.17.1 or later.
- The I/O module that failed must be a cluster and HA I/O module in slot 4 and you must be replacing it with an equivalent cluster and HA I/O module. You cannot change the I/O module type.
- Your storage system configuration must have only one cluster and HA I/O module located in slot 4, not two cluster and HA I/O modules.
- Your storage system must be a two-node (switchless or switched) cluster configuration.
- The controller with the failed cluster and HA I/O module (the impaired controller) must have already taken over the healthy partner controller. The takeover should have occurred automatically if the I/O module is failed.

  For two-node clusters, the storage system cannot discern which controller has the failed I/O module, so either controller might initiate the takeover. The cluster and HA I/O module hot-swap procedure is only supported when the controller with the failed I/O module (the impaired controller) has taken over the healthy controller.

You can verify that the impaired controller successfully took over the healthy controller by entering the `storage failover show` command.

If you are not sure which controller has the failed I/O module, contact NetApp Support.

- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Step 2: Prepare the storage system and I/O module slot 4

Prepare the storage system and I/O module slot 4 so that it is safe to remove the failed cluster and HA I/O module:

Steps

1. Properly ground yourself.
2. Unplug cabling from the failed cluster and HA I/O module.

   Make sure to label the cables so that later in this procedure you can reconnect them to the same ports.

3. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<number of hours down>h`

   For example, the following AutoSupport message suppresses automatic case creation for two hours:

   `node2::> system node autosupport invoke -node * -type all -message MAINT=2h`

4. Disable automatic giveback:
   a. Enter the following command from the console of the impaired controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
5. Prepare the failed cluster and HA module in slot 4 for removal by removing it from service and powering it off:
   a. Enter the following command:

      `system controller slot module remove -node impaired_node_name -slot slot_number`

   b. Enter `y` when you see the prompt *Do you want to continue?*

      For example, the following command prepares the module in slot 4 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

```
node2::> system controller slot module remove -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered off for removal.

Do you want to continue? {y|n}: y

The module has been successfully removed from service and powered
off. It can now be safely removed.
```

6. Verify the failed cluster and HA module in slot 4 is powered off:

   `system controller slot module show`

   The output should show *`powered-off`* in the status column for the failed module in slot 4.

Step 3: Replace the failed cluster and HA I/O module

Replace the failed cluster and HA I/O module in slot 4 with an equivalent I/O module:

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the failed cluster and HA I/O module from the impaired controller:

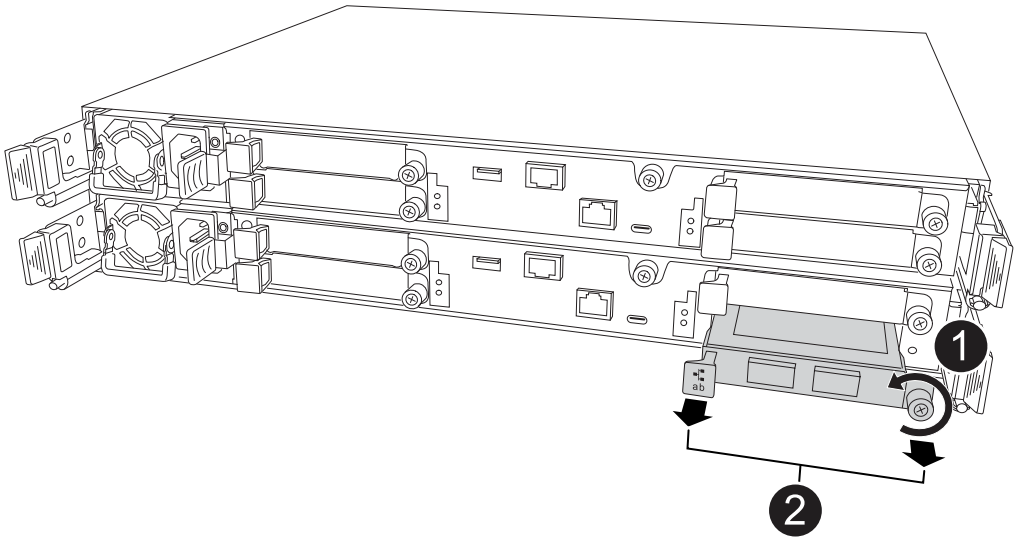

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew on the right.

3. Install the replacement cluster and HA I/O module into slot 4:

a. Align the I/O module with the edges of the slot.
b. Gently push the I/O module all the way into the slot, making sure to properly seat the I/O module into the connector.

   You can use the tab on the left and the thumbscrew on the right to push in the I/O module.

c. Turn the thumbscrew clockwise to tighten.

4. Cable the cluster and HA I/O module.

Step 4: Bring the replacement cluster and HA I/O module online

Bring the replacement cluster and HA I/O module in slot 4 online, verify the module ports initialized successfully, verify slot 4 is powered on, and then verify the module is online and recognized.

Steps

1. Bring the replacement cluster and HA I/O module online:
   a. Enter the following command:

      `system controller slot module insert -node impaired_node_name -slot slot_name`

   b. Enter `y` when you see the prompt, *Do you want to continue?*

      The output should confirm the cluster and HA I/O module was successfully brought online (powered on, initialized, and placed into service).

      For example, the following command brings slot 4 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered on and initialized.

Do you want to continue? {y|n}: `y`

The module has been successfully powered on, initialized and placed
into service.
```

2. Verify that each port on the cluster and HA I/O module successfully initialized:

   `event log show -event *hotplug.init*`

   > It might take several minutes to allow for any required firmware updates and port initialization.

   The output should show a hotplug.init.success EMS event logged for each port on the cluster and HA I/O module with *`hotplug.init.success:`* in the *`Event`* column.

For example, the following output shows initialization succeeded for cluster and HA I/O module ports e4b and e4a:

```
node2::> event log show -event *hotplug.init*

Time                Node             Severity      Event

------------------- ---------------- -------------
---------------------------

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4b" in slot 4 succeeded

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4a" in slot 4 succeeded

2 entries were displayed.
```

3. Verify I/O module slot 4 is powered on and ready for operation:

   `system controller slot module show`

   The output should show slot 4 status as *`powered-on`* and therefore ready for operation of the replacement cluster and HA I/O module.

4. Verify that the replacement cluster and HA I/O module is online and recognized.

   Enter the command from the console of the impaired controller:

   `system controller config show -node local -slot4`

   If the replacement cluster and HA I/O module was successfully brought online and is recognize, the output shows I/O module information, including port information, for slot 4.

   For example, you should see output similar to the following:

```
node2::> system controller config show -node local -slot 4

Node: node2
Sub- Device/
Slot slot Information
---- ---- ------------------------------
   4    - Dual 40G/100G Ethernet Controller CX6-DX
                  e4a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2807GJFM-B
                  e4b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2809G26F-A
                  Device Type:        CX6-DX PSID(NAP0000000027)
                  Firmware Version:   22.44.1700
                  Part Number:        111-05341
                  Hardware Revision:  20
                  Serial Number:      032403001370
```

Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the healthy controller, restoring automatic giveback, and reenabling AutoSupport automatic case creation.

Steps

1. Return the healthy controller (the controller that was taken over) to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode healthy_node_name
   ```

2. Restore automatic giveback from the console of the impaired controller (the controller that took over the healthy controller):

   ```
   storage failover modify -node local -auto-giveback true
   ```

3. If AutoSupport is enabled, restore automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=end
   ```

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an I/O module in your FAS50 storage system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

Before you begin

All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace a failed I/O module

To replace a failed I/O module, locate it in the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug cabling from the failed I/O module.

   Make sure to label the cables so that you know where they came from.

3. Remove the failed I/O module from the controller:

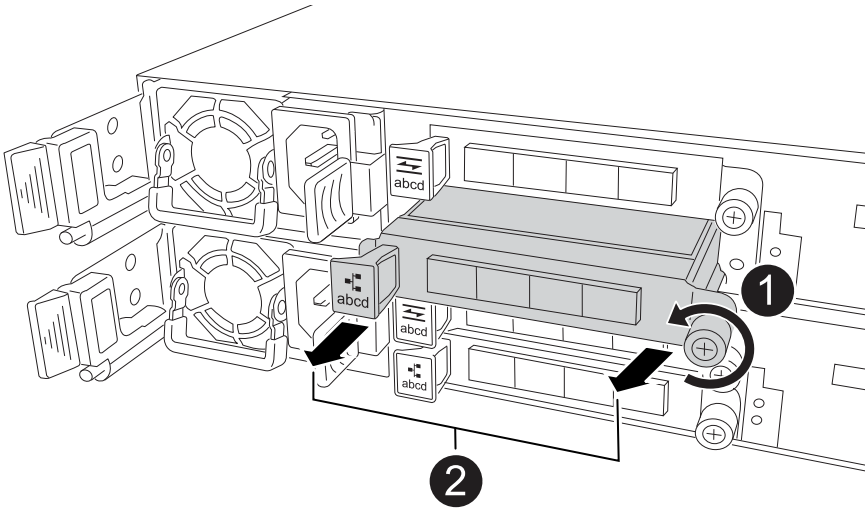

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the replacement I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

c. Turn the thumbscrew clockwise to tighten.

5. Cable the I/O module.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```

3. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```

4. If AutoSupport is enabled, restore automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NV battery - FAS50

Replace the NV battery in your FAS50 storage system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

Before you begin

All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

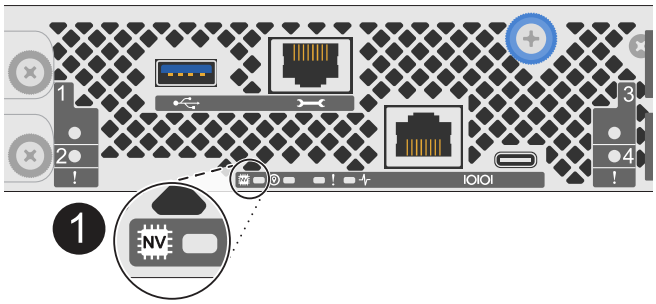

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

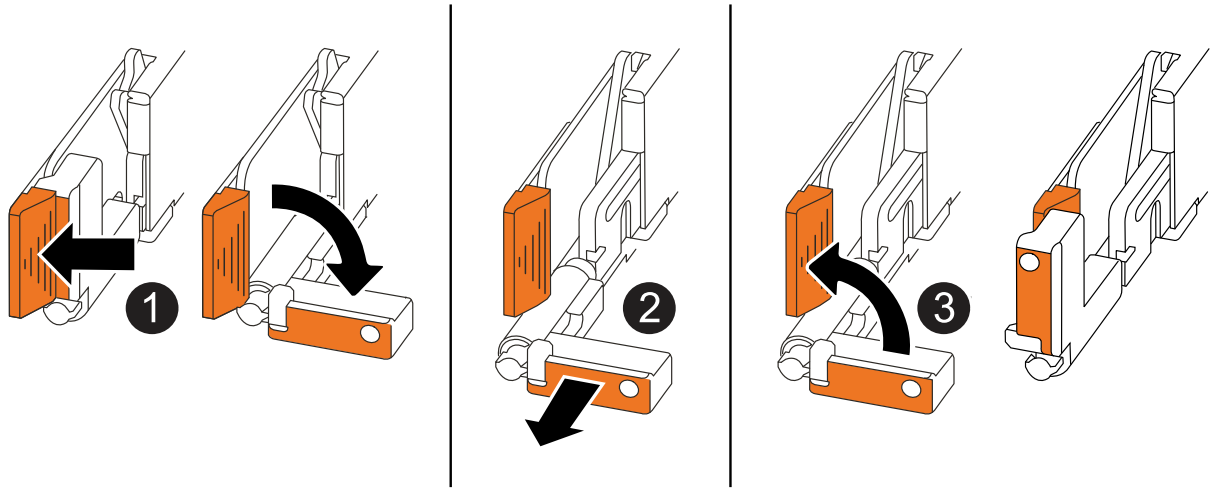

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the NV battery

Remove the failed NV battery from the controller and install the replacement NV battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the NV battery.
3. Remove the NV battery:

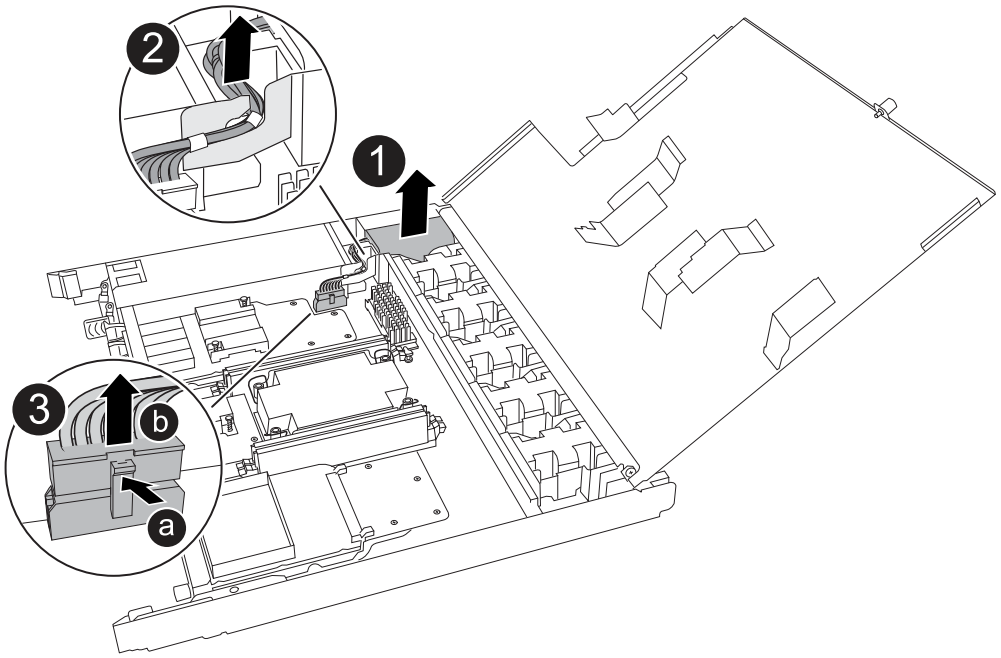

1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

4. Install the replacement NV battery:
   a. Remove the replacement battery from its package.
   b. Plug the wiring connector into its socket.
   c. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   d. Place the NV battery into its compartment.

      The NV battery should sit flush in its compartment.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

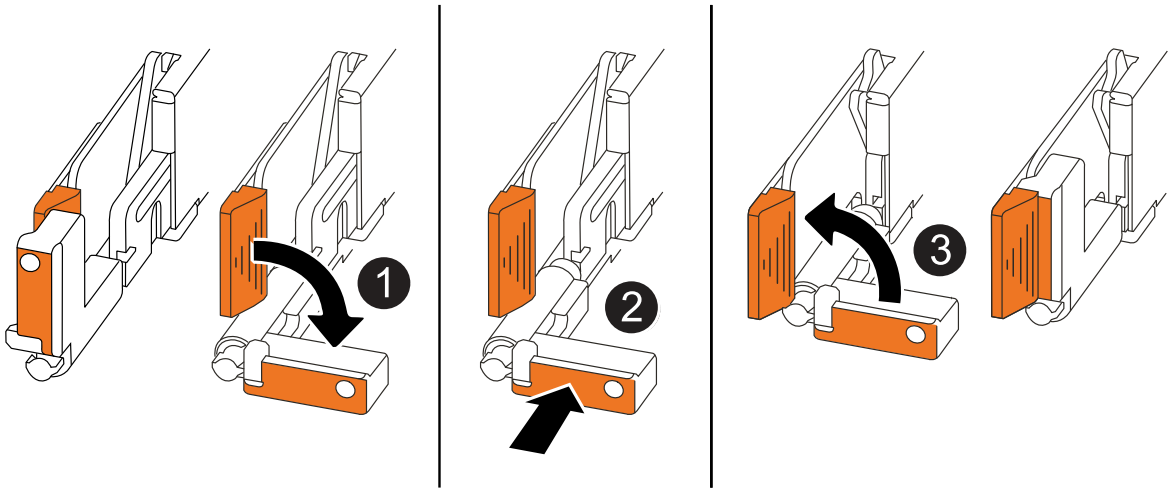

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.

The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```

8. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - FAS50

Replace an AC or DC power supply unit (PSU) in your FAS50 storage system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cord, replacing the faulty PSU, and then reconnecting it to the power source.

About this task

- This procedure is written for replacing one PSU at a time.

The PSUs are redundant and hot-swappable.

- IMPORTANT: Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.
- Use the appropriate procedure for your type of PSU: AC or DC.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the power cord from the PSU by opening the power cord retainer, and then unplug the power cord from the PSU.

   PSUs do not have a power switch.

4. Remove the PSU:

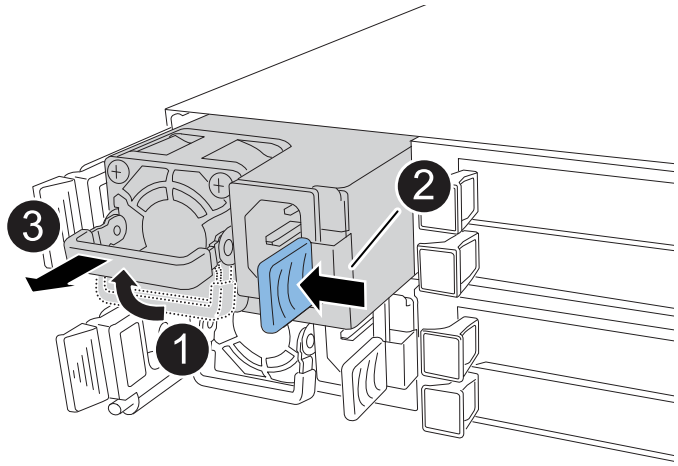

1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

5. Install the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

c. Rotate the handle down, so it is out of the way of normal operations.

6. Reconnect the power cord to the PSU and secure the power cord with the power cord retainer.

   Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the PSU:

   PSUs do not have a power switch.

   a. Unscrew the two thumb screws on the D-SUB DC power cord connector.

      The illustration and table in step 4 shows the two thumb screws (item #1) and the D-SUB DC power cord connector (item #2).

   b. Unplug the cord from the PSU and set it aside.

4. Remove the PSU:

   a. Rotate the handle up, to its horizontal position, and then grasp it.

   b. With your thumb, press the terracotta tab to release the locking mechanism.

   c. Pull the PSU out of the controller while using your other hand to support its weight.

   The PSU is short. Always use two hands to support it when removing it from the controller so that it does not swing free from the controller and injure you.

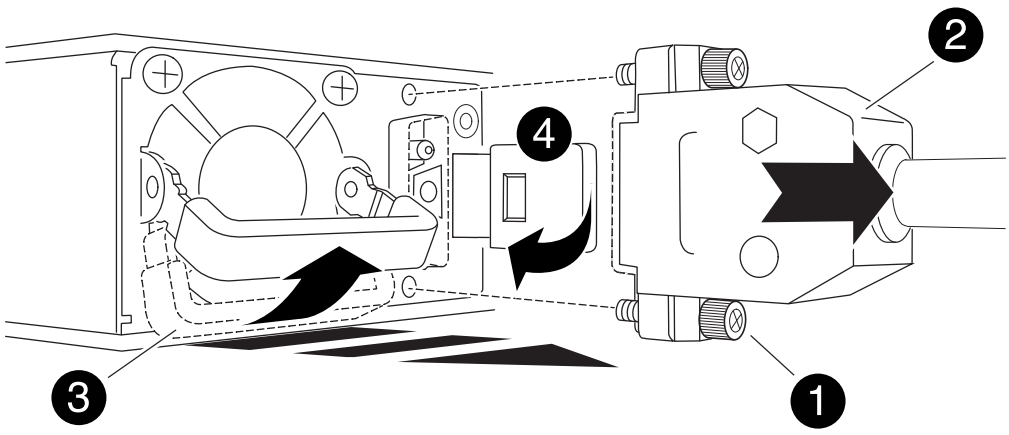

1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

5. Insert the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently slide the PSU into the controller until the locking tab clicks into place.

      A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.
6. Reconnect the D-SUB DC power cord:

   Once power is restored to the PSU, the status LED should be green.

   a. Plug the D-SUB DC power cord connector into the PSU.
   b. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - FAS50

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your FAS50 storage system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

About this task

- You can use this procedure with all versions of ONTAP supported by your storage system.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  A storage system has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led *off*` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

You must remove the controller from the chassis when you replace the controller or replace a component inside the controller.

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.

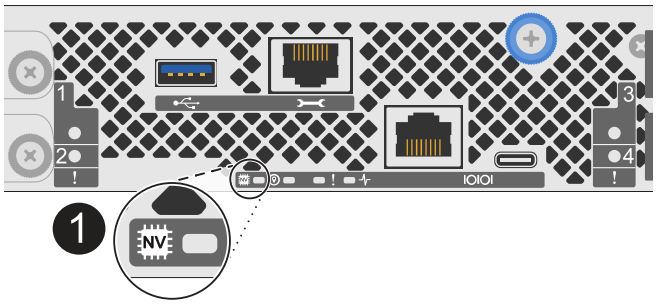

1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:

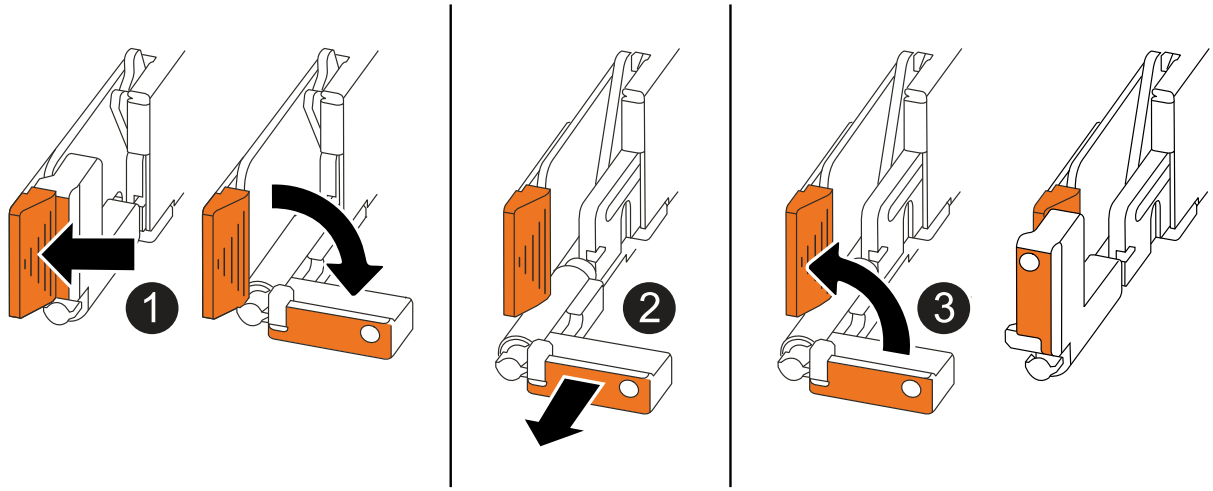

1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the RTC battery

Remove the failed RTC battery and install the replacement RTC battery.

Steps

1. Locate the RTC battery.
2. Remove the RTC battery:

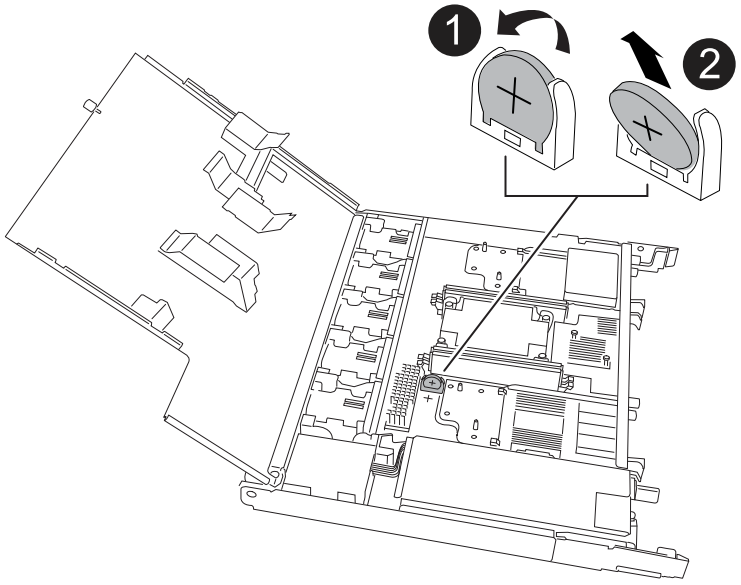

1	Gently rotate the RTC battery at an angle away from its holder.
2	Lift the RTC battery out of its holder.

3. Install the replacement RTC battery:
   a. Remove the replacement battery from the antistatic shipping bag.
   b. Position the battery so that the plus sign on the battery faces out to correspond with the plus sign on the motherboard.
   c. Insert the battery into the holder at an angle, and then push it into an upright position so it is fully seated in the holder.
   d. Visually inspect the battery to make sure that it is completely seated in its holder and that the polarity is correct.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.

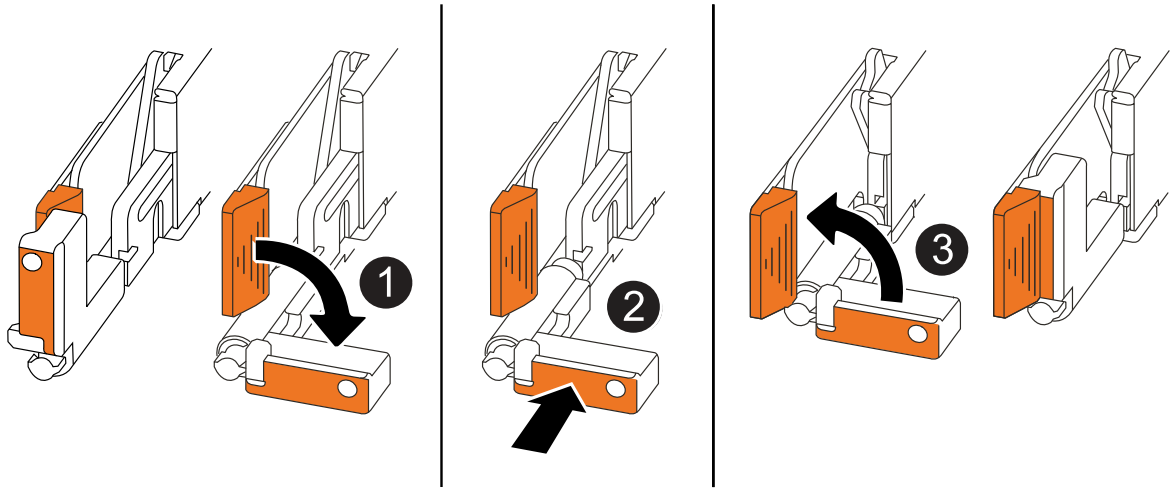

1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   Do not completely insert the controller in the chassis until instructed to do so.

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

   Once power is restored to the PSU, the status LED should be green.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode impaired_node_name`

8. Restore automatic giveback from the console of the healthy controller:

   `storage failover modify -node local -auto-giveback true`

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

   `system node autosupport invoke -node * -type all -message MAINT=END`

Step 5: Reset the time and date on the controller

> After replacing the RTC battery, inserting the controller and powering on first BIOS reset, you will see the following error messages:
> `RTC date/time error. Reset date/time to default`
> `RTC power failure error`
> These messages are expected and you can continue with this procedure.

1. On the healthy controller, check the date and time:

   `cluster date show`

   > If your storage system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*.

2. On the impaired controller, at the LOADER prompt, check the time and date:

   `cluster date show`

   a. If necessary, modify the date:

      `set date mm/dd/yyyy`

   b. If necessary, set the time, in GMT:

```
set time hh:mm:ss
```

c. Confirm the date and time.

3. At the LOADER prompt, enter `bye` to reinitialize the I/O modules, other components, and let the controller reboot.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

FAS70 and FAS90 systems

Install and setup

Installation and configuration workflow - FAS70 and FAS90

To install and configure your FAS70 or FAS90 system, you review the hardware requirements, prepare your site, install and cable the hardware components, power on the system, and set up your ONTAP cluster.

Review installation requirements

Review the equipment and tools needed to install your storage system and storage shelves and review the lifting and safety precautions.

Prepare to install the FAS70 or FAS90 storage system

To prepare to install your system, you need to get the site ready, check the environmental and electrical requirements, and ensure there's enough rack space. Then, unpack the equipment, compare its contents to the packing slip, and register the hardware to access support benefits.

Install the hardware for the FAS70 or FAS90 storage system

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system in the cabinet or telco rack. Next, slide the shelves onto the rails. Finally, attach cable management devices to the rear of the storage system for organized cable routing.

Cable the controllers and storage shelves for the FAS70 or FAS90 storage system

To cable the hardware, first connect the storage controllers to your network and then connect the controllers to your storage shelves.

Power on the FAS70 or FAS90 storage system

Before you power on the controllers, power on each NS224 shelf and assign a unique shelf ID to ensure each shelf is uniquely identified within the setup, connect the laptop or console to the controller, and then connect the controllers to the power sources.

6 Set up your cluster

After you've powered on your storage system, you set up your cluster.

Installation requirements - FAS70 and FAS90

Review the equipment needed and the lifting precautions for your FAS70 or FAS90 storage system and storage shelves.

Equipment needed for install

To install your storage system, you need the following equipment and tools.

- Access to a Web browser to configure your storage system
- Electrostatic discharge (ESD) strap
- Flashlight
- Laptop or console with a USB/serial connection
- Phillips #2 screwdriver

Lifting precautions

Storage systems and shelves are heavy. Exercise caution when lifting and moving these items.

Storage system weight

Take the necessary precautions when moving or lifting your storage system.

An A1K storage system can weigh up to 62.83 lbs (28.5 kg). To lift the storage system, use two people or a hydraulic lift.

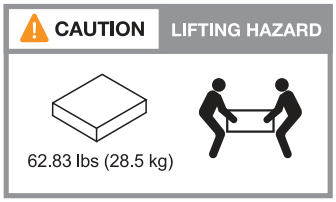

A FAS70 or FAS90 storage system can weigh up to 62.83 lbs (28.5 kg). To lift the system, use two people or a hydraulic lift.

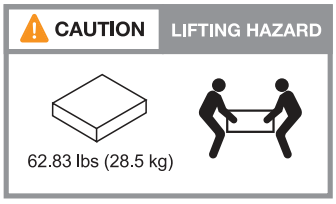

Shelf weight

Take the necessary precautions when moving or lifting your shelf.

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep

all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

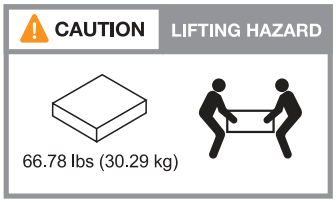

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.

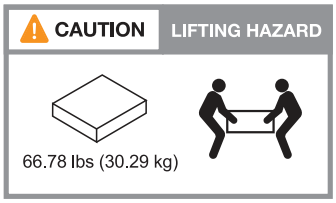

A DS460C shelf can weigh up to 260.4 lbs (181.1 kg). To lift the storage shelf, you might need up to five people or a hydraulic lift. Keep all components in the storage shelf (both front and rear) to prevent unbalancing the shelf weight.

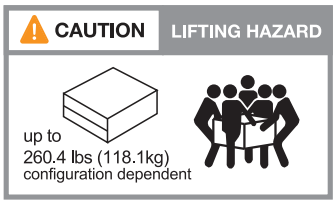

Related information

- Safety information and regulatory notices

What's next?

After you've reviewed the hardware requirements, you prepare to install your FAS70 or FAS90 storage system.

Prepare to install - FAS70 and FAS90

Prepare to install your FAS70 or FAS90 storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the system to access support benefits.

Step 1: Prepare the site

To install your storage system, ensure that the site and the cabinet or rack that you plan to use meet specifications for your configuration.

Steps

1. Use NetApp Hardware Universe to confirm that your site meets the environmental and electrical requirements for your storage system.
2. Make sure you have adequate cabinet or rack space for your storage system, shelves, and any switches:
   - 4U in an HA configuration
   - 2U for each NS224 storage shelf

3. Install any required network switches.

   See the Switch documentation for installation instructions and NetApp Hardware Universe for compatibility information.

Step 2: Unpack the boxes

After you've ensured that the site and the cabinet or rack that you plan to use for your storage system meet the required specifications, unpack all boxes and compare the contents to the items on the packing slip.

Steps

1. Carefully open all the boxes and lay out the contents in an organized manner.
2. Compare the contents you've unpacked with the list on the packing slip.

   You can get your packing list by scanning the QR code on the side of the shipping carton.

   The following items are some of the contents you might see in the boxes.

   Ensure that everything in the boxes matches the list on the packing slip. If there are any discrepancies, note them down for further action.

Hardware

- Bezel
- Cable management device
- Storage system
- Rail kits with instructions (optional)
- Storage shelf (if you ordered additional storage)

Cables

- Management Ethernet cables (RJ-45 cables)
- Network cables
- Power cords
- Storage cables (if you ordered additional storage)
- USB-C serial console cable

Step 3: Register your storage system

After you've ensured that your site meets the requirements for your storage system specifications, and you've verified that you have all the parts you ordered, you should register your storage system.

Steps

1. Locate the System Serial Numbers (SSN) for every controller being installed. You can find the serial numbers in the following locations:
2. You can find the serial numbers in the following locations:
   - On the packing slip
   - In your confirmation email
   - On each controller's System Management module

3. Go to the NetApp Support Site.
4. Determine whether you need to register your storage system:

If you are a…	Follow these steps…
Existing NetApp customer	a. Sign in with your username and password. b. Select **Systems** > **My Systems**. c. Confirm that the new serial numbers are listed. d. If it is not, follow the instructions for new NetApp customers.
New NetApp customer	a. Click **Register Now**, and create an account. b. Select **Systems** > **Register Systems**. c. Enter the storage system's serial numbers and requested details. After your registration is approved, you can download any required software. The approval process might take up to 24 hours.

What's next?

After you've prepared to install your FAS70 or FAS90 hardware, you install the hardware for your FAS70 or FAS90 storage system.

Install the hardware - FAS70 and FAS90

After you prepare to install your FAS70 or FAS90 storage system, install the hardware for the system. First, install the rail kits. Then install and secure your platform in a cabinet or telco rack.

Skip this step if your cabinet is pre-populated.

Before you begin

- Make sure you have the instructions packaged with the rail kit.
- Be aware of the safety concerns associated with the weight of the storage system and shelf.
- Understand that the airflow through the storage system enters from the front where the bezel or end caps are installed and exhausts out the rear where the ports are located.

Steps

1. Install the rail kits for your storage system and shelves as needed, using the instructions included with the kits.
2. Install and secure your storage system in the cabinet or telco rack:
   a. Position the storage system onto the rails in the middle of the cabinet or telco rack, and then support the storage system from the bottom and slide it into place.
   b. Make sure that the guiding pins of the cabinet or telco rack are securley in the chassis guide slots.
   c. Secure the storage system to the cabinet or telco rack using the included mounting screws.
3. Attach the bezel to the front of the storage system.
4. Attach the cable management devices to the rear of the storage system.

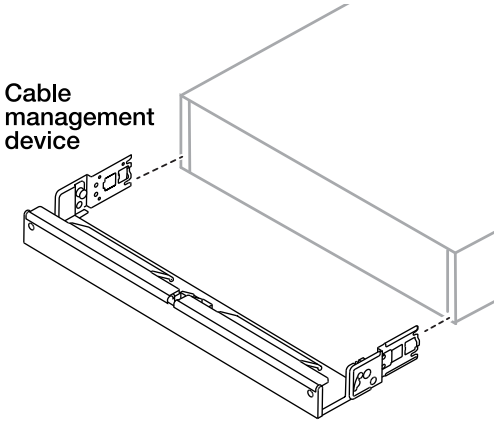

5. Install and secure the shelf as needed.
   a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

      If you are installing multiple shelves, place the first shelf directly above the controllers. Place the second shelf directly under the controllers. Repeat this pattern for any additional shelves.

   b. Secure the shelf to the cabinet or telco rack using the included mounting screws.

      What's next?

      After you've installed the hardware for your FAS70 or FAS90 storage system, you cable the hardware for your FAS70 or FAS90 storage system..

Cable the hardware - FAS70 and FAS90

After you install the rack hardware for your FAS70 or FAS90 storage system, install the network cables for the controllers, and connect the cables between the controllers and storage shelves.

Before you begin

Contact your network administrator for information about connecting the storage system to the switches.

About this task

- These procedures show common configurations. The specific cabling depends on the components ordered for your storage system. For comprehensive configuration and slot priority details, see NetApp Hardware Universe.
- The I/O slots on FAS70 and FAS90 controllers are numbered 1 through 11.

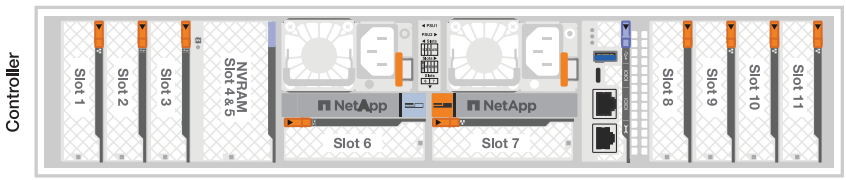

- The cabling graphics have arrow icons showing the proper orientation (up or down) of the cable connector pull-tab when inserting a connector into a port.

  As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it

over and try again.

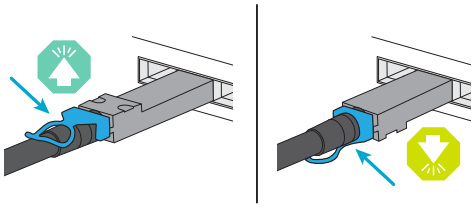

- If cabling to an optical switch, insert the optical transceiver into the controller port before cabling to the switch port.

Step 1: Connect the storage controllers to your network

Cable the controllers to your ONTAP cluster. This procedure differs depending on your storage system model and I/O module configuration.

The cluster interconnect traffic and the HA traffic share the same physical ports.

Switchless cluster cabling

Use the Cluster/HA interconnect cable to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A to port e1a on Controller B.
2. Connect port e7a on Controller A to port e7a on Controller B.

Cluster/HA interconnect cables

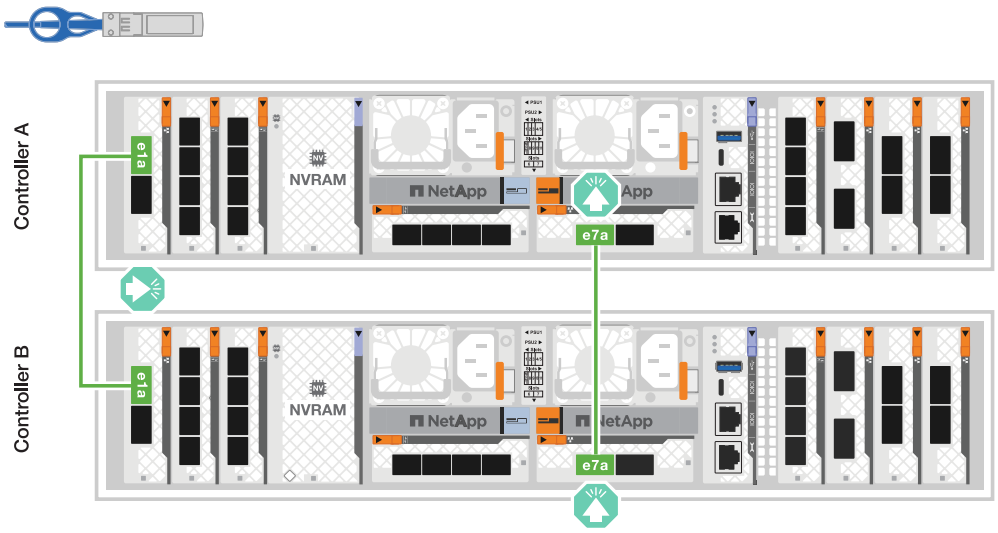

Switched cluster cabling

Use the 100 GbE cable to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A and port e1a on Controller B to cluster network switch A.
2. Connect port e7a on Controller A and port e7a on Controller B to cluster network switch B.

100 GbE cable

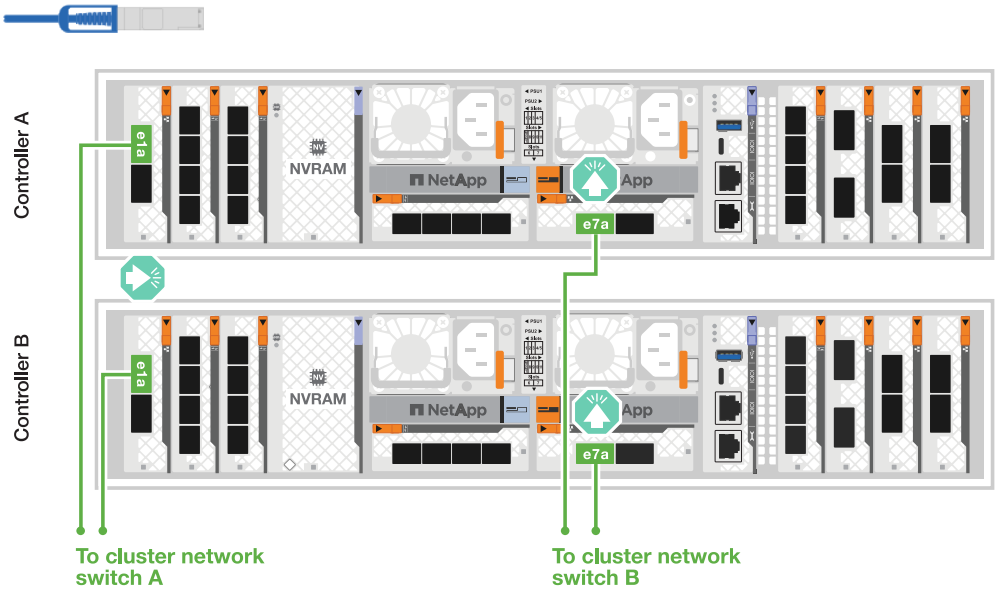

Step 2: Cable the host network connections

Connect the Ethernet module ports to your host network.

The following are some typical host network cabling examples. See NetApp Hardware Universe for your specific system configuration.

Steps

1. Connect ports e9a and e9b to your Ethernet data network switch.

   For maximum system performance for cluster and HA traffic, do not use ports e1b and e7b ports for host network connections. Use a separate host card to maximize performance.

   100 GbE cable

   

   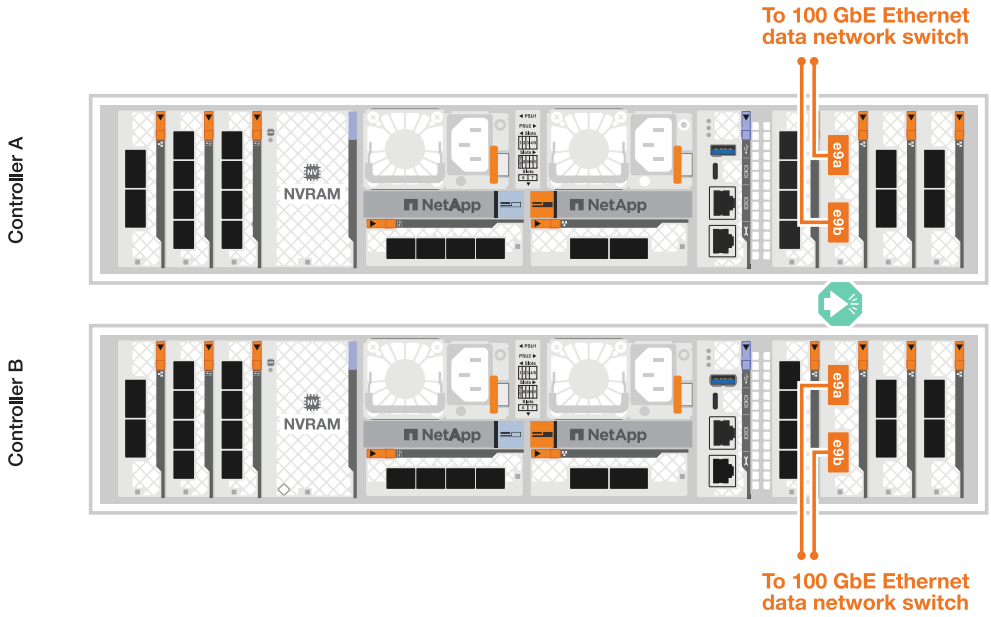

2. Connect your 10/25 GbE host network switches.

   4-ports, 10/25 GbE Host

   

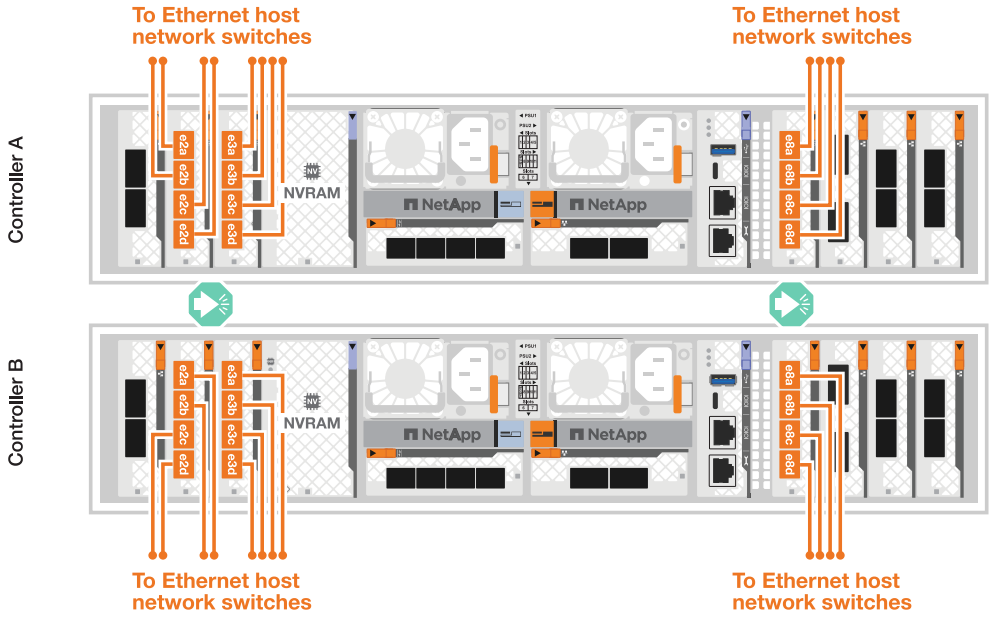

Step 3: Cable the management network connections

Use the 1000BASE-T RJ-45 cables to connect the management (wrench) ports on each controller to the management network switches.

1000BASE-T RJ-45 cables

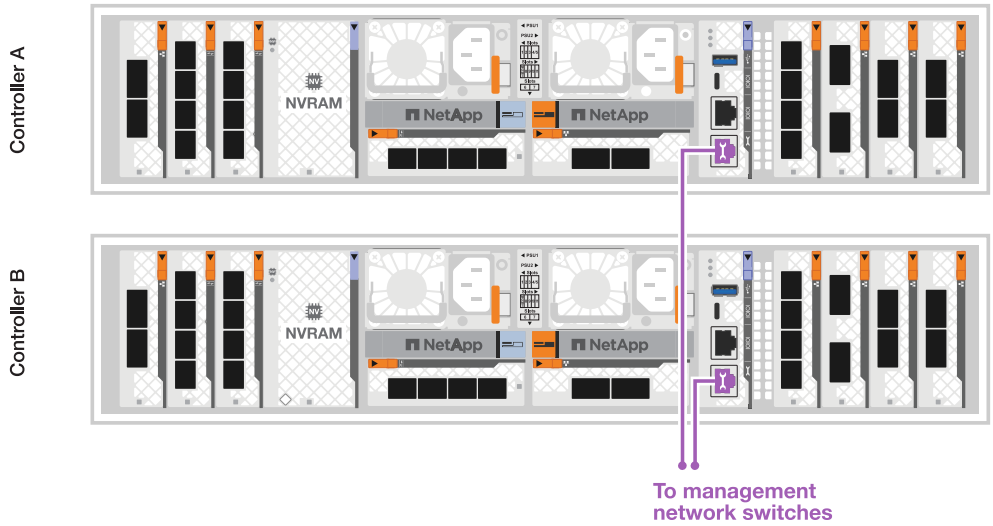

Do not plug in the power cords yet.

Step 4: Cable the shelf connections

The following cabling procedures show how to connect your controllers to a storage shelf. Choose one of the following cabling options that matches your setup.

For the maximum number of shelves supported for your storage system and for all of your cabling options, see NetApp Hardware Universe.

For additional SAS shelf cabling guidance, see SAS cabling rules and concepts - shelves with IOM12/IOM12B modules.

About this task

The FAS70 and FAS90 storage systems supports DS212C, DS224C, DS460C, and NS224 shelves with either the NSM100 or NSM100B module.

The major differences between the NS224 modules are:

- NSM100 shelf modules use built-in ports e0a and e0b.
- NSM100B shelf modules use ports e1a and e1b in slot 1.

The following NS224 cabling example shows NSM100 modules in the NS224 shelves when referring to shelf module ports.

Option 1: One NS224 storage shelf

Connect each controller to the NSM modules on the NS224 shelf. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables

Steps

1. On controller A, connect the following ports:
   a. Connect port e11a to NSM A port e0a.
   b. Connect port e11b to port NSM B port e0b.

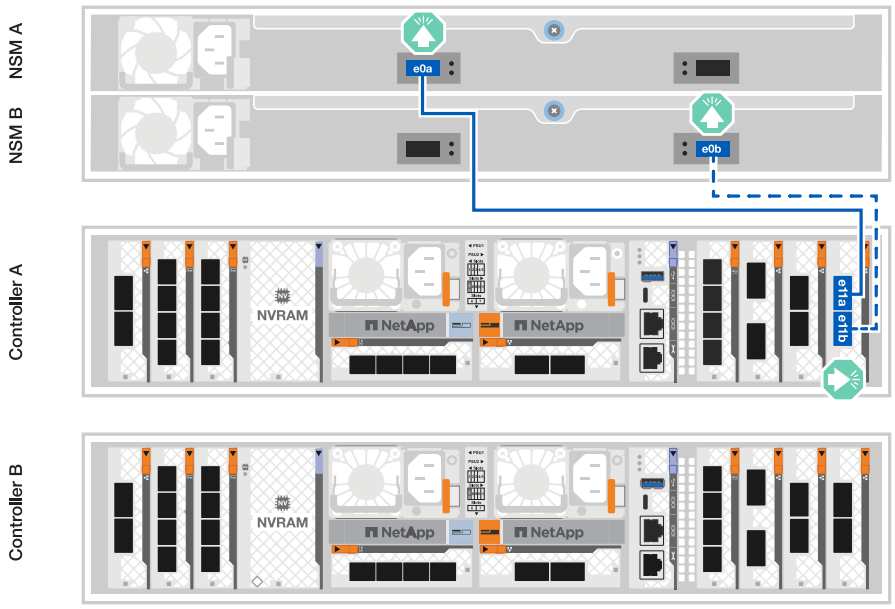

2. On controller B, connect the following ports:
   a. Connect port e11a to NSM B port e0a.
   b. Connect port e11b to NSM A port e0b.

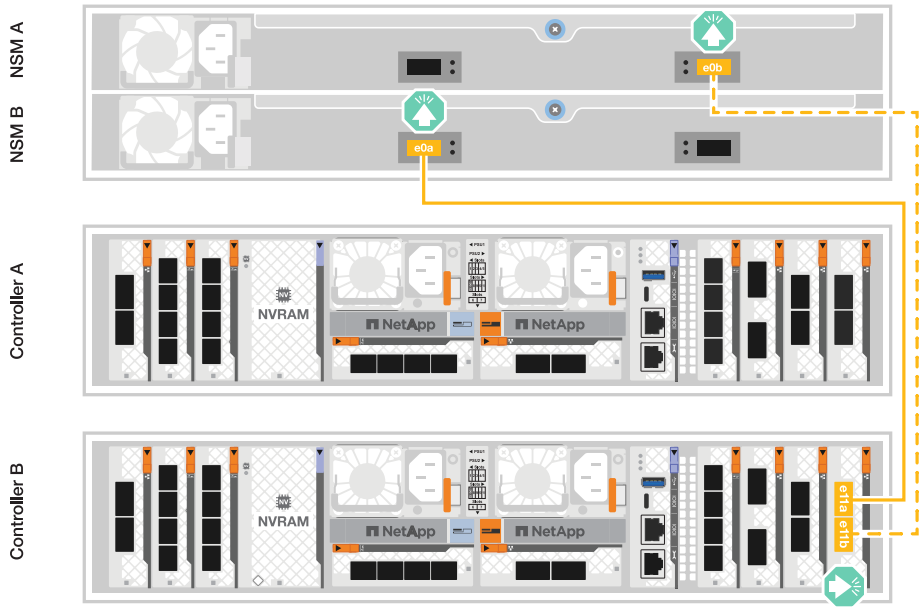

Option 2: Two NS224 storage shelves

Cable each controller to the NSM modules on both NS224 shelves. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables

Steps

1. On controller A, connect the following ports:

    a. Connect port e11a to shelf 1 NSM A port e0a.

    b. Connect port e11b to shelf 2 NSM B port e0b.

    c. Connect port e10a to shelf 2 NSM A port e0a.

    d. Connect port e10b to shelf 1 NSM A port e0b.

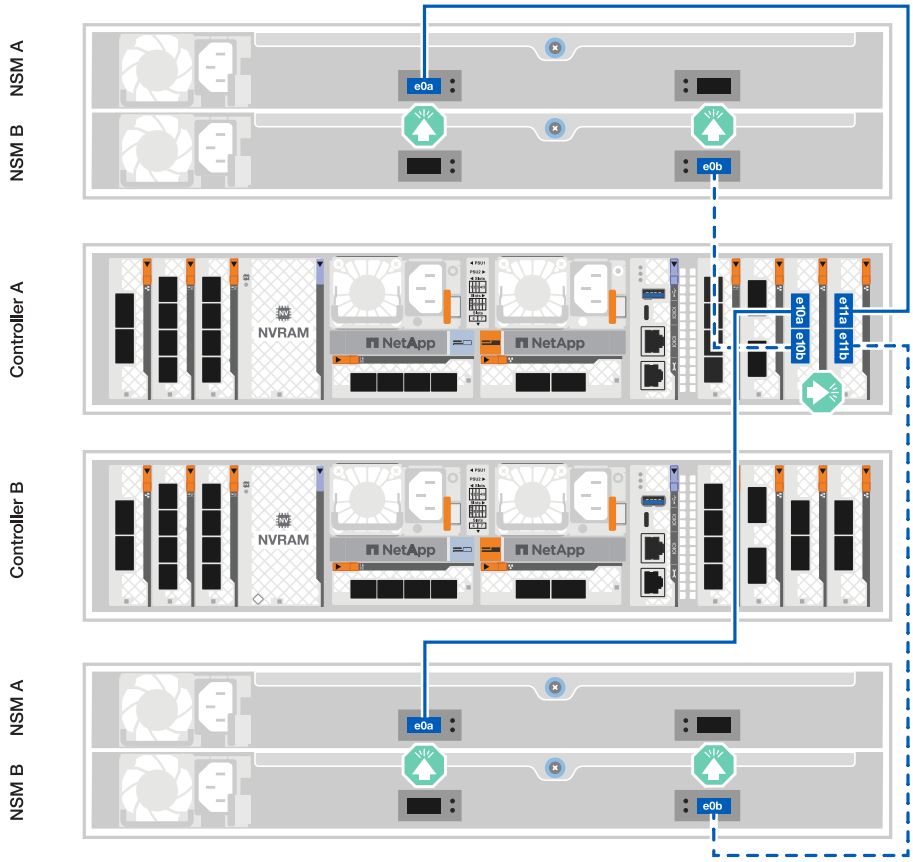

2. On controller B, connect the following ports:

   a. Connect port e11a to shelf 1 NSM B port e0a.
   b. Connect port e11b to shelf 2 NSM A port e0b.
   c. Connect port e10a to shelf 2 NSM B port e0a.
   d. Connect port e10b to shelf 1 NSM A port e0b.

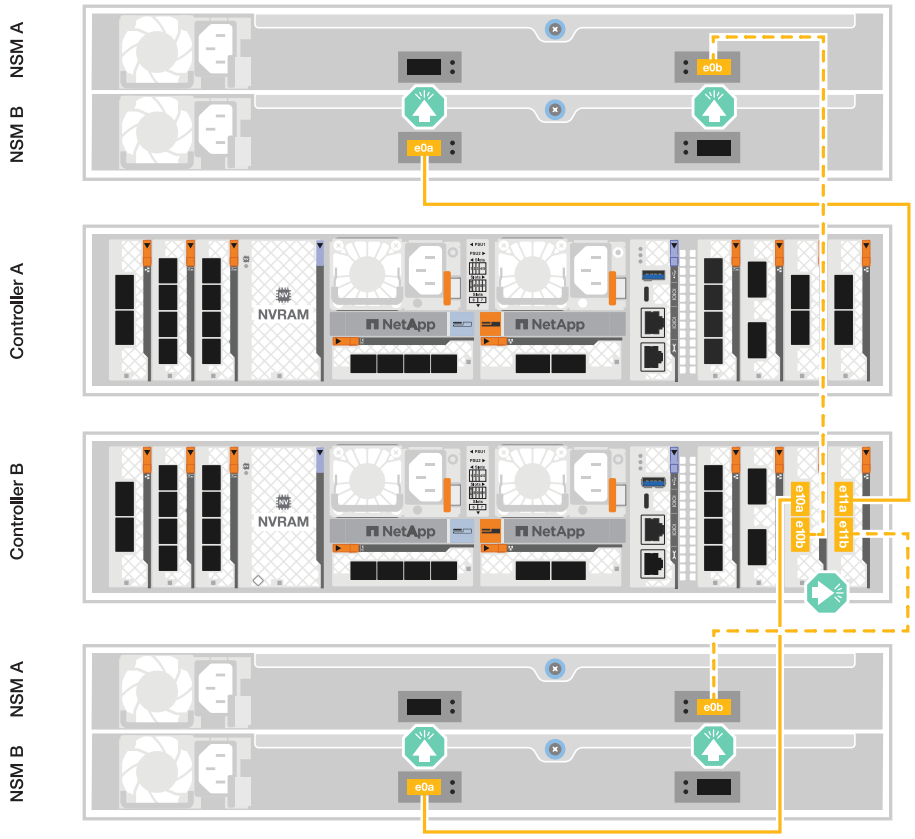

Option 3: Two DS460C shelves

Cable each controller to the IOM modules on both DS460C shelves. The graphics show controller A cabling in blue and controller B cabling in yellow.

mini-SAS HD cable

Steps

1. On controller A, cable the following connections:

    a. Connect port e10a to shelf 1 IOM A port 1.

    b. Connect port e10c to shelf 2 IOM A port 1

    c. Connect port e11b to shelf 1 IOM B port 3.

    d. Connect port e11d to shelf 2 IOM B port 3.

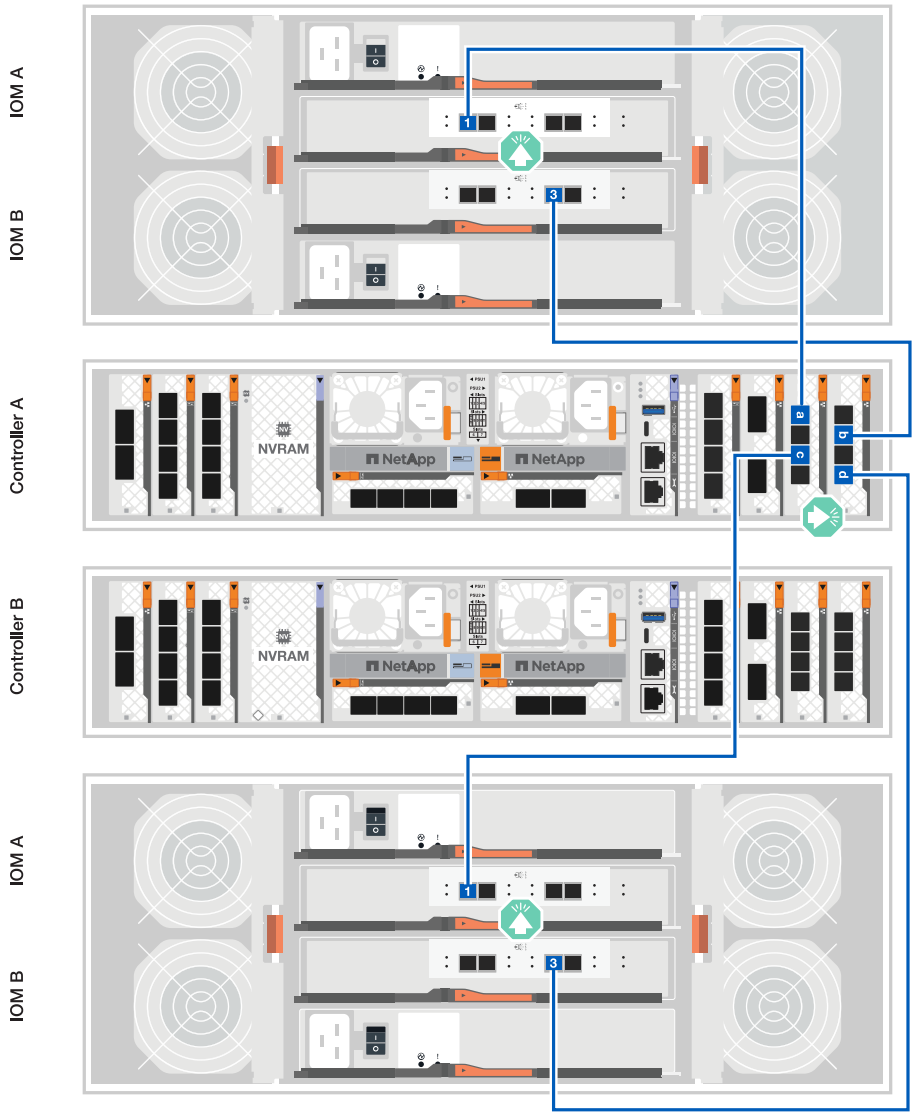

2. On controller B, cable the following connections:

   a. Connect port e10a to shelf 1 IOM B port 1.
   b. Connect port e10c to shelf 2 IOM B port 1.
   c. Connect port e11b to shelf 1 IOM A port 3.
   d. Connect port e11d to shelf 2 IOM A port 3.

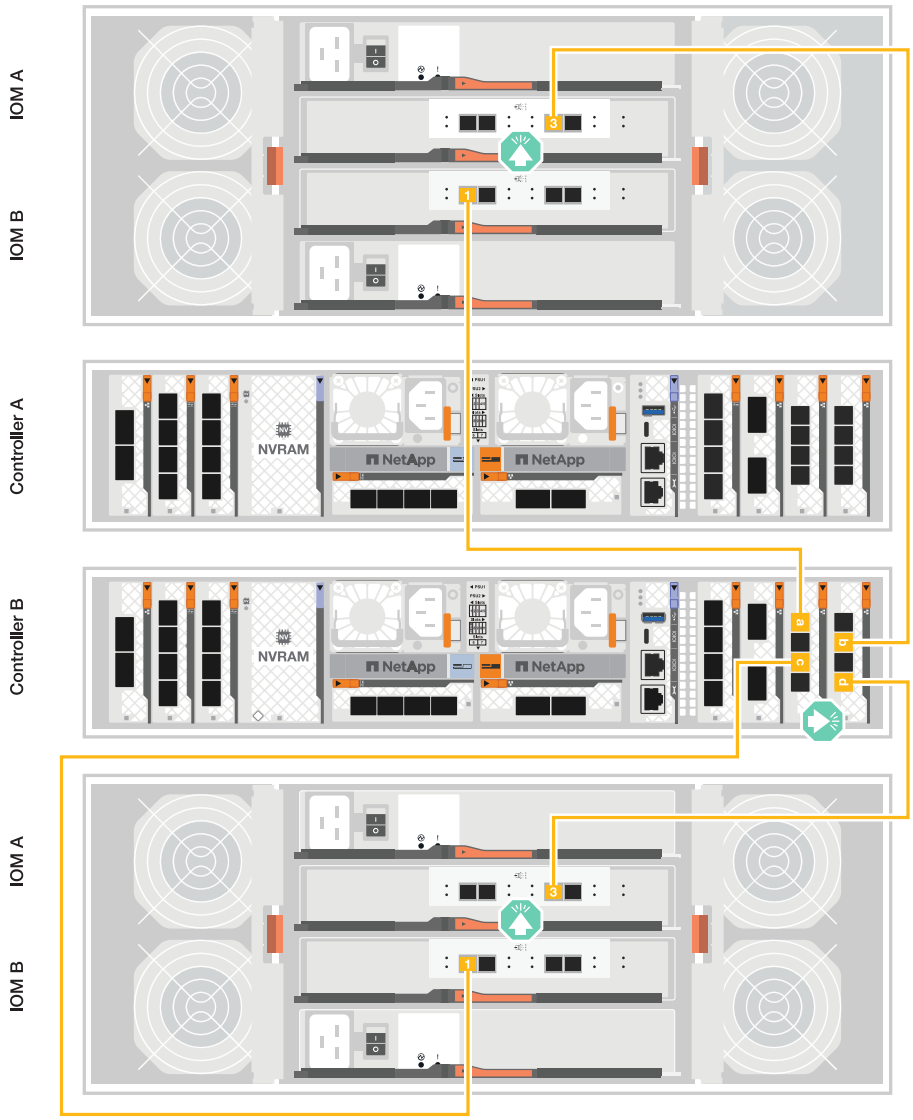

What's next?

After you've cabled the hardware for your FAS70 or FAS90 system, you power on the FAS70 or FAS90 storage system.

Power on the storage system - FAS70 and FAS90

After you install the rack hardware for your FAS70 0r FAS90 storage system and install the cables for the controllers and storage shelves, you should power on your storage shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

Option 1: NS224 shelves

Each shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup.

Before you begin

Make sure you have a paperclip or narrow tipped ball point pen for setting NS224 storage shelf IDs.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed shelf ID of 00.

- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, and then plug them back in) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, and then connecting the power cords to power sources on different circuits.

   The shelf powers on and boots automatically when plugged into the power source.

2. Remove the left end cap to access the shelf ID button behind the faceplate.

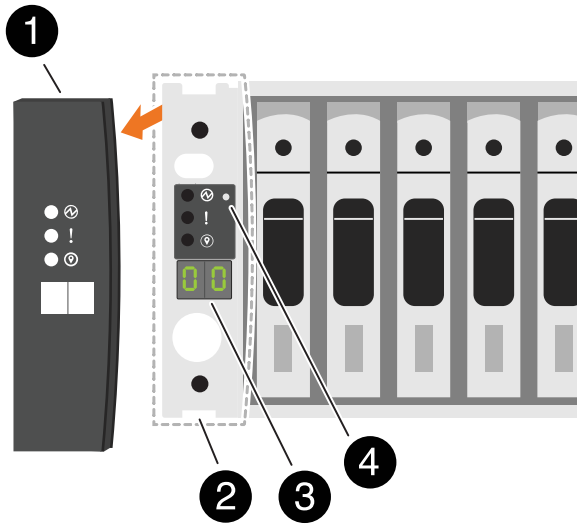

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID number
4	Shelf ID button

3. Change the first number of the shelf ID:

a. Insert the straightened end of a paperclip or narrow tipped ball point pen into the small hole to press the shelf ID button.
b. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

   It can take up to 15 seconds for the number to blink. This activates the shelf ID programming mode.

   > If the ID takes longer than 15 seconds to blink, press and hold the shelf ID button again, making sure to press it in all the way.

c. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

   Each press and release duration can be as short as one second.

   The first number continues to blink.

4. Change the second number of the shelf ID:
   a. Press and hold the button until the second number on the digital display blinks.

      It can take up to three seconds for the number to blink.

      The first number on the digital display stops blinking.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

      The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the shelf ID button until the second number stops blinking.

   It can take up to three seconds for the number to stop blinking.

   Both numbers on the digital display start blinking and the amber LED illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.
   a. Unplug the power cord from both power supplies on the shelf.
   b. Wait 10 seconds.
   c. Plug the power cords back into the shelf power supplies to complete the power cycle.

      A power supply is powered on as soon as the power cord is plugged in. Its bicolored LED should illuminate green.

7. Replace the left end cap.

Option 2: SAS shelves

Each shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed shelf ID of 00.

- You must power cycle a shelf (turn off the power switch on each of the power supplies of the SAS shelf, wait the appropriate amount of time, and then switch the power back on) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, connecting the power cords to power sources on different circuits, and then turning on the power switch on each of the power supplies (at the rear of the shelf).

   The shelf powers on and boots automatically when powered on.

2. Remove the left end cap to access the orange shelf ID button on the faceplate.

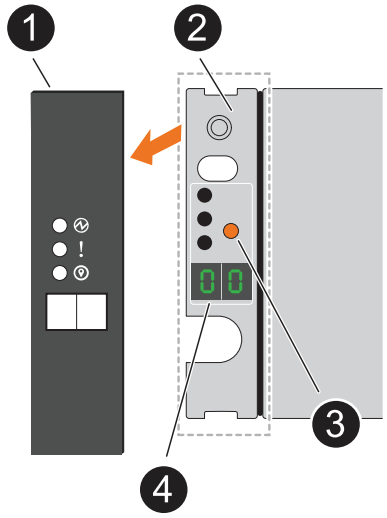

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID button
4	Shelf ID number

3. Change the first number of the shelf ID:
   a. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

      It can take up to 15 seconds for the number to blink. This activates the shelf ID programming mode.

If the ID takes longer than 15 seconds to blink, press and hold the shelf ID button again, making sure to press it in all the way.

b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

   Each press and release duration can be as short as one second.

   The first number continues to blink.

4. Change the second number of the shelf ID:
   a. Press and hold the button until the second number on the digital display blinks.

      It can take up to three seconds for the number to blink.

      The first number on the digital display stops blinking.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

      The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the shelf ID button until the second number stops blinking.

   It can take up to three seconds for the number to stop blinking.

   Both numbers on the digital display start blinking and the amber LED illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.
   a. Turn off the power switch on each of the power supplies.
   b. Wait 10 seconds.
   c. Turn on the power switch on each of the power supplies to complete the power cycle.

      When a power supply is powered on, the bicolored LED should illuminate green.

7. Replace the left end cap.

Step 2: Power on the controllers

After you've powered on your shelves and assigned them unique IDs, power on the storage controllers.

Steps

1. Connect your laptop to the serial console port. This will allow you to monitor the boot sequence when the controllers are powered on.
   a. Set the serial console port on the laptop to 115,200 baud with N-8-1.

      See your laptop's online help for instructions on how to configure the serial console port.

   b. Using the console cable provided with your storage system, connect one end of the console cable to

your laptop and the other end to the serial console port on controller A.

c. Connect the laptop to the switch on the management subnet.

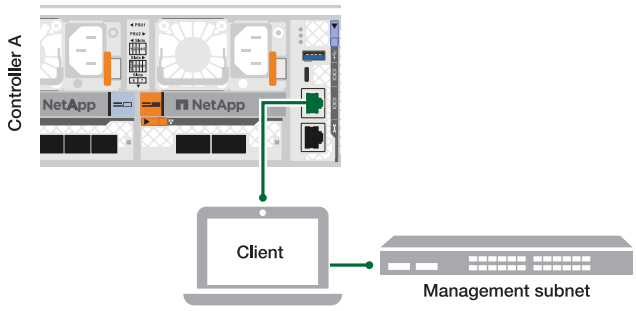

2. Assign a TCP/IP address to the laptop, using one that is on the management subnet.
3. Plug the two power cords into the controller power supplies, and then connect them to power sources on different circuits.

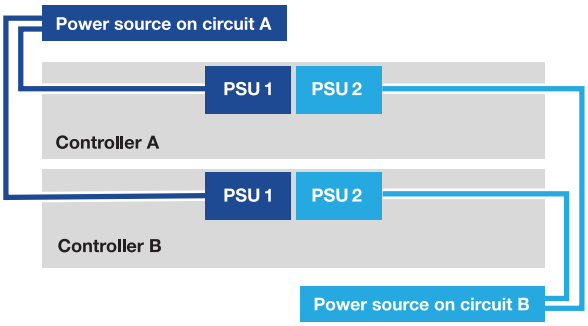

- The system begins to boot. Initial booting might take up to eight minutes.
- The LEDs flash on and the fans start, which indicates that the controllers are powering on.
- The fans might be very noisy when they first start up. The fan noise during start-up is normal.

4. Secure the power cords using the securing device on each power supply.

What's next?

After you've turned on your FAS70 or FAS90 storage system, you set up your cluster.

Maintain

Overview of the maintenance procedures - FAS70 and FAS90

Maintain the hardware of your FAS70 or FAS90 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS70 or FAS90 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS70 and FAS90 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media- manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot the image from a USB drive and restore the configuration from the partner node
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Fan	A fan cools the controller.
Flash Cache	Flash Cache speeds access to data through real-time intelligent caching of recently read user data and NetApp metadata. It's effective for random read-intensive workloads, including databases, email, and file services.
NVRAM	The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module.
NV battery	The NV battery is responsible for providing power to the NVRAM module while data in-flight is being destaged to flash memory after a power loss.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.
System management module	The System management module provides the interface between the controller and a console or laptop for controller or system maintenance purposes. The System management module contains the boot media and stores the system serial number (SSN).

Boot media - automated recovery

Boot media automated recovery workflow - FAS70 and FAS90

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your FAS70 or FAS90 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the System Management module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - FAS70 and FAS90

Before replacing the boot media in your FAS70 or FAS90 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming the cluster ports on the impaired controller are working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Review the following requirements.

- You must replace the failed boot media with a replacement boot media you received from NetApp.
- The cluster ports are used to communicate between the two controllers during the automated boot recovery process. Make sure that the cluster ports on the impaired controller are working properly.
- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg
  - /cfcard/kmip/certs/client.crt
  - /cfcard/kmip/certs/client.key
  - /cfcard/kmip/certs/CA.pem
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - FAS70 and FAS90

Shut down the impaired controller in your FAS70 or FAS90 storage system to prevent data loss and maintain system stability during the automatic boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - FAS70 and FAS90

The boot media in your FAS70 or FAS90 storage system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media in the System Management module, and then reinstalling the System Management module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the System Management module and is accessed by removing the module from the system.

Replace the boot media.

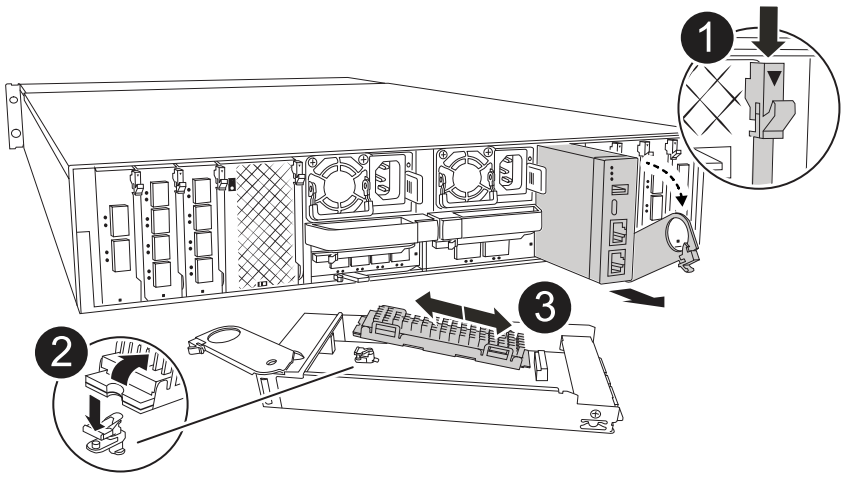

1	System Management module cam latch
2	Boot media locking button
3	Boot media

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.

   If your storage system has DC power supplies, disconnect the power cable block from the power supply units (PSUs).

3. Remove the System Management module:
   a. Remove any cables connected to the System Management module. Make sure that you label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   c. Depress the System Management cam button.
   d. Rotate the cam latch down as far as it will go.
   e. Remove the System Management module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
   f. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:
   a. Press the blue locking button.
   b. Rotate the boot media up, slide it out of the socket, and set it aside.
5. Install the replacement boot media into the System Management module:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down toward the locking button.

c. Push the locking button, rotate the boot media all the way down and then release the locking button.

6. Reinstall the System Management module:

   a. Align the module with the edges of the enclosure slot opening.

   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.

7. Rotate the cable management tray up to the closed position.

   a. Recable the System Management module.

8. Plug the power cables into the power supplies and reinstall the power cable retainer.

   The controller begins to boot as soon as power is reconnected to the system.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - FAS70 and FAS90

After installing the new boot media device in your FAS70 or FAS90 storage system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media part to NetApp - FAS70 and FAS90

If a component in your FAS70 or FAS90 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - FAS70 and FAS90

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the FAS70 or FAS90 system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

If your system is running in ONTAP 9.17.1 and later, use the automatic boot recovery procedure.

To get started, review the recovery requirements, shut down the controller, replace the boot media, use the USB drive to restore the image, and reapply encryption settings if necessary.

Review boot media replacement requirements

Review the requirements for replacing the boot media.

Check onboard encryption keys

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the impaired controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONTAP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - FAS70 and FAS90

Before replacing the boot media in your FAS70 or FAS90 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption support for manual boot media recovery - FAS70 and FAS90

To ensure data security on your FAS70 or FAS90 storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

What's next?

After checking the encryption key support and status on the boot media, you need to shut down the controller.

Shut down the controller for manual boot media recovery - FAS70 and FAS90

Shut down the impaired controller in your FAS70 or FAS90 storage system to prevent data loss and maintain system stability during the manual boot media recovery process.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After shutting down the controller, you need to replace the boot media.

Replace the boot media and prepare for manual boot recovery - FAS70 and FAS90

The boot media in your FAS70 or FAS90 system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

Step 1: Replace the boot media

The boot media is located inside the System Management module and is accessed by removing the module from the system.

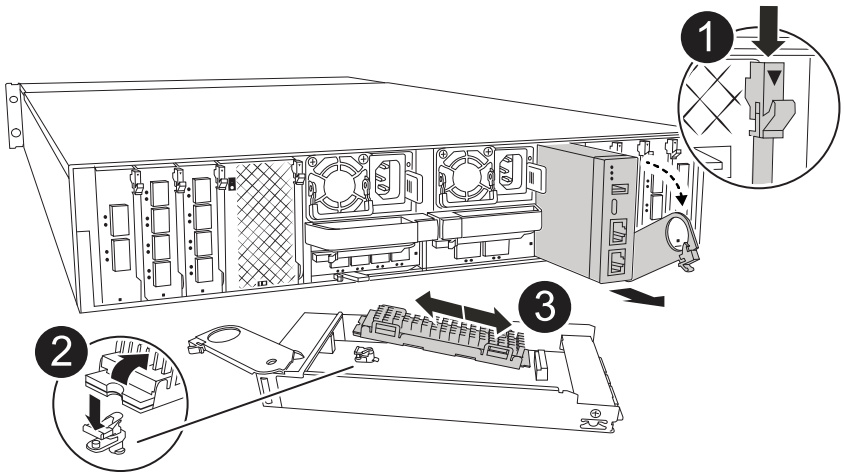

1	System Management module cam latch
2	Boot media locking button
3	Boot media

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.

   If your storage system has DC power supplies, disconnect the power cable block from the power supply units (PSUs).

3. Remove the System Management module:
   a. Remove any cables connected to the System Management module. Make sure that you label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   c. Depress the System Management cam button.
   d. Rotate the cam latch down as far as it will go.
   e. Remove the System Management module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
   f. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:
   a. Press the blue locking button.
   b. Rotate the boot media up, slide it out of the socket, and set it aside.
5. Install the replacement boot media into the System Management module:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the locking button.
   c. Push the locking button, rotate the boot media all the way down and then release the locking button.
6. Reinstall the System Management module.
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
7. Rotate the cable management tray up to the closed position.
   a. Recable the System Management module.

Step 2: Transfer the ONTAP image to the boot media

The replacement boot media that you installed is without an ONTAP image. You can transfer the ONTAP image to the replacement boot media by downloading the appropriate ONTAP service image from the NetApp Support Site to a USB flash drive and then to the replacement boot media.

Before you begin

- You must have an empty USB flash drive, formatted to FAT32, with at least 4GB capacity.
- Download a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site. Use the `version -v` command to display if your version of ONTAP supports NVE. If the command output displays

`<10no- DARE>`, your version of ONTAP does not support NVE.

  - If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image from the Downloads link on the page, to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      The USB flash drive should have the appropriate ONTAP image of what the impaired controller is running.

   c. Remove the USB flash drive from your laptop.
2. Insert the USB flash drive into the USB slot on the System Management module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

3. Plug the power cables back into the PSU.

   The system will begin to reboot, typically to the LOADER prompt.

4. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

What's next?

After replacing the boot media, you need to boot the recovery image.

Manual boot media recovery from a USB drive - FAS70 and FAS90

After installing the new boot media device in your FAS70 or FAS90 system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.

3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.

5. Give back the controller using the `storage failover giveback -fromnode local` command.

6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

   NOTE: If the process fails, contact NetApp Support.

What's next?

After booting the recovery image, you need to restore encryption on the boot media.

Restore encryption keys after manual boot recovery - FAS70 and FAS90

Restore encryption on the replacement boot media in your FAS70 or FAS90 system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*********************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*********************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

What's next?

After restoring encryption on the boot media, you need to return the failed part to NetApp.

Return the failed part to NetApp - FAS70 and FAS90

If a component in your FAS70 or FAS90 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Controller

Controller replacement workflow - FAS70 and FAS90

Get started with replacing the controller in your FAS70 or FAS90 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

1 Review controller replacement requirements

To replace the controller module, you must meet certain requirements.

2 Shut down the impaired controller

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

3 Replace the controller

Replacing the controller includes removing the impaired controller, moving the FRU components to the replacement controller module, and then installing the replacement controller module in the enclosure.

4 Restore and verify the system configuration

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

5 Give back the controller

Transfer the ownership of storage resources back to the replacement controller.

6 Complete controller replacement

Verify the LIfs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - FAS70 and FAS90

Before replacing the controller in your FAS70 or FAS90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements for replacing the controller.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- Do not use this procedure for controller upgrades; instead, refer to the Choose your controller hardware upgrade procedure for guidance.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this controller replacement procedure.
- You must replace the failed component with the field-replaceable unit (FRU) you received from NetApp.
- You must replace a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.

- You cannot change any drives or drive shelves as part of this procedure.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text log file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

What's next?

After you've reviewed the requirements to replace your FAS70 or FAS90 controller, you need to shut down the impaired controller.

Shut down the impaired controller - FAS70 and FAS90

Shut down the controller in your FAS70 or FAS90 storage system to prevent data loss and ensure system stability when replacing the controller.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After you've shut down the controller, you need to replace the controller.

Replace the controller - FAS70 and FAS90

Replace the controller in your FAS70 or FAS90 system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

Step 1: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

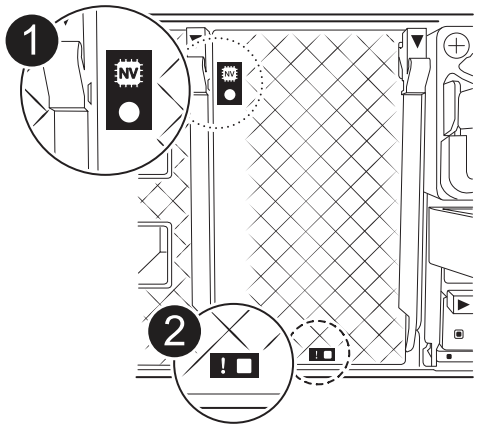

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

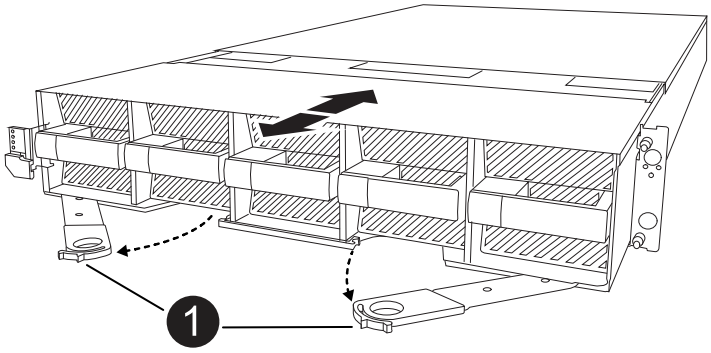

1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 2: Move the fans

You must remove the five fan modules from the impaired controller module to the replacement controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press the gray locking button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

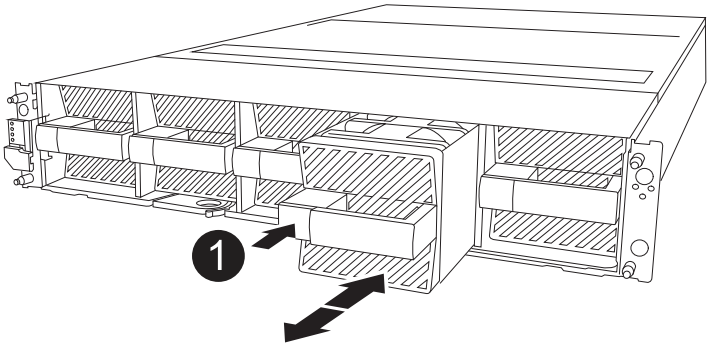

1	Black locking button

4. Install the fan in the replacement controller module:
   a. Align the edges of the fan housing with the opening in the front of the replacement controller module.
   b. Gently slide the fan module all the way into the replacement controller module until it locks in place.
5. Repeat the preceding steps for the remaining fan modules.

Step 3: Move the NV battery

Move the NV battery to the replacement controller.

Steps

1. Open the NV battery air duct cover and locate the NV battery.

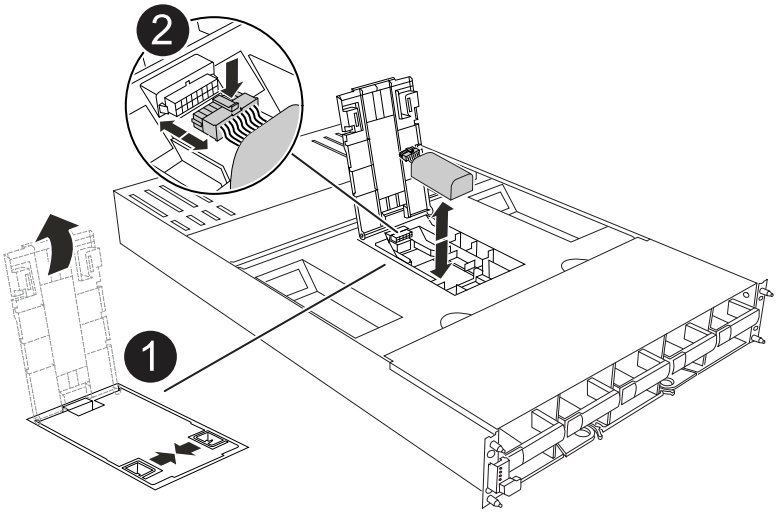

1	NV battery air duct cover
2	NV battery plug
3	NV battery pack

2. Lift the battery up to access the battery plug.

3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module.
5. Move the battery pack to the replacement controller module and then install it in the NV battery air duct:
   a. Open the NV battery air duct in the replacement controller module.
   b. Plug the battery plug into the socket and make sure that the plug locks into place.
   c. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   d. Close the air duct cover.

Step 4: Move system DIMMs

Move the DIMMs to the replacement controller module.

Steps

1. Open the motherboard air duct and locate the DIMMs.

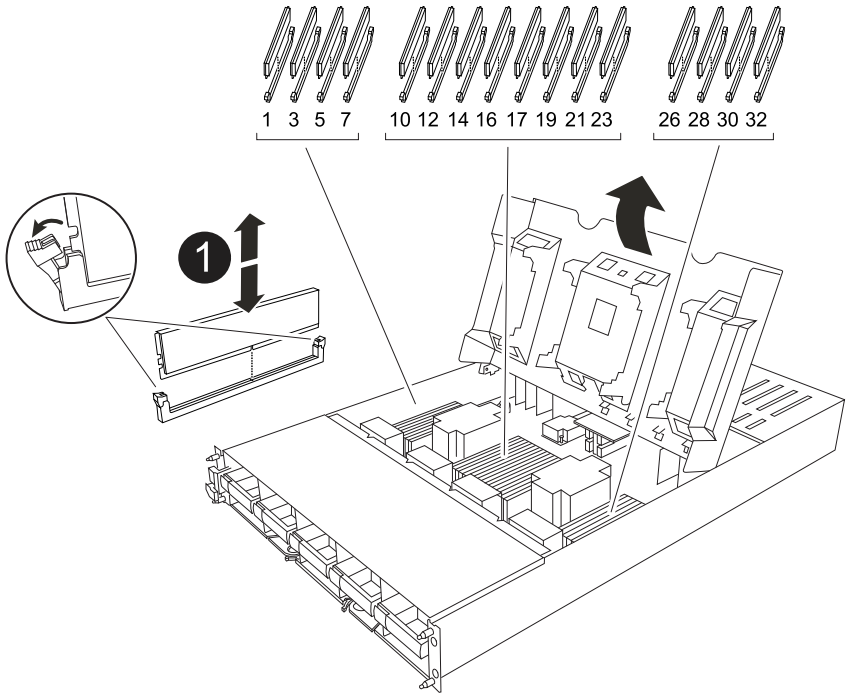

1	System DIMM

2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

> Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

4. Locate the slot where you are installing the DIMM in the replacement controller module.
5. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Repeat these steps for the remaining DIMMs.
   Close the motherboard air duct.

Step 5: Install the controller module

Reinstall the controller module and boot it.

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

   Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

   The controller boots to the LOADER prompt as soon as it is fully seated.

4. From the LOADER prompt, enter `show date` to display the date and time on the replacement controller. Date and time are in GMT.

   Time displayed is local time not always GMT and is displayed in 24hr mode.

5. Set the current time in GMT with the `set time hh:mm:ss` command. You can get the current GMT from the partner node the `date -u`command.
6. Recable the storage system, as needed.

   If you removed the transceivers (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

What's next?

After you've replaced the impaired FAS70 or FAS90 controller, you need to restore the system configuration.

Restore and verify the system configuration - FAS70 and FAS90

Verify that the controller's HA configuration is active and functioning correctly in your FAS70 or FAS90 storage system, and confirm that the system's adapters list all the paths

to the disks.

Step 1: Verify HA config settings

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

Steps

1. Boot to maintenance mode: `boot_ontap maint`
   a. Enter `y` when you see *Continue with boot?*.

      If you see the *System ID mismatch* warning message, enter `y`.

2. Enter `sysconfig -v` and capture the display contents.

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

5. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha`

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed: `ha-config show`

Step 2: Verify disk list

Steps

1. Verify that the adapter lists the paths to all disks with the `storage show disk -p`.

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode: `halt`.

What's next?

After you've restored and verified the system configuration for your FAS70 or FAS90 system, you need to give back the controller.

Give back the controller - FAS70 and FAS90

Return control of storage resources to the replacement controller so your FAS70 or FAS90 system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`
   - If you encounter errors, contact NetApp Support.
9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

    c. Verify status of the keys: `security key-manager key query -restored false`

The output should show no results when when properly synchronized.

If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

After you've transferred the ownership of storage resources back to the replacement controller, you need to complete the controller replacement procedure.

Complete controller replacement - FAS70 and FAS90

To complete the controller replacement for your AFF A1K system, first restore the NetApp Storage Encryption configuration (if necessary). Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, return the failed part to NetApp.

Step 1: Verify LIFs and check cluster health

Before returning the replacement node to service, verify that the logical interfaces are on their home ports, check the cluster health, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any logical interfaces are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS70 and FAS90

Replace a DIMM in your FAS70 or FAS90 system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system

from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- Make sure all other components in the system are functioning properly; if not, you must contact technical support.
- Make sure you replace the failed component with a replacement component you received from NetApp.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

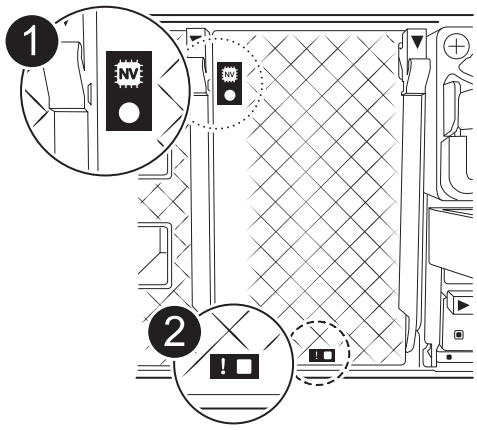

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

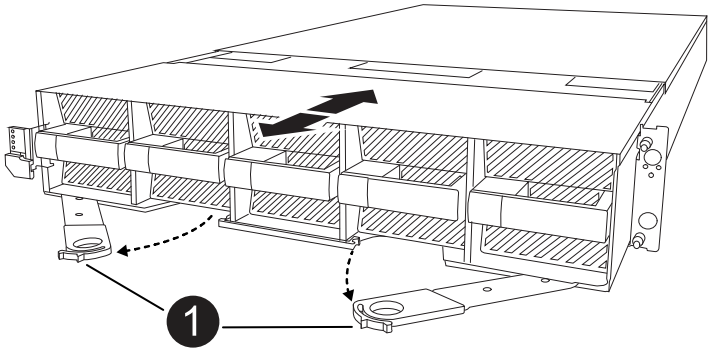

1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace a DIMM

You must replace a DIMM when the system reports a permanent failure condition for that DIMM.

1. If you are not already grounded, properly ground yourself.
2. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
3. Locate the DIMMs on your controller module and identify the DIMM for replacement.

   DIMMs locations are dependent on the system model:

Model	DIMM slot location
FAS70	Slots 3, 10, 19, 26
FAS90	Slots 3, 7, 10, 14, 19, 23, 26, 30

4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

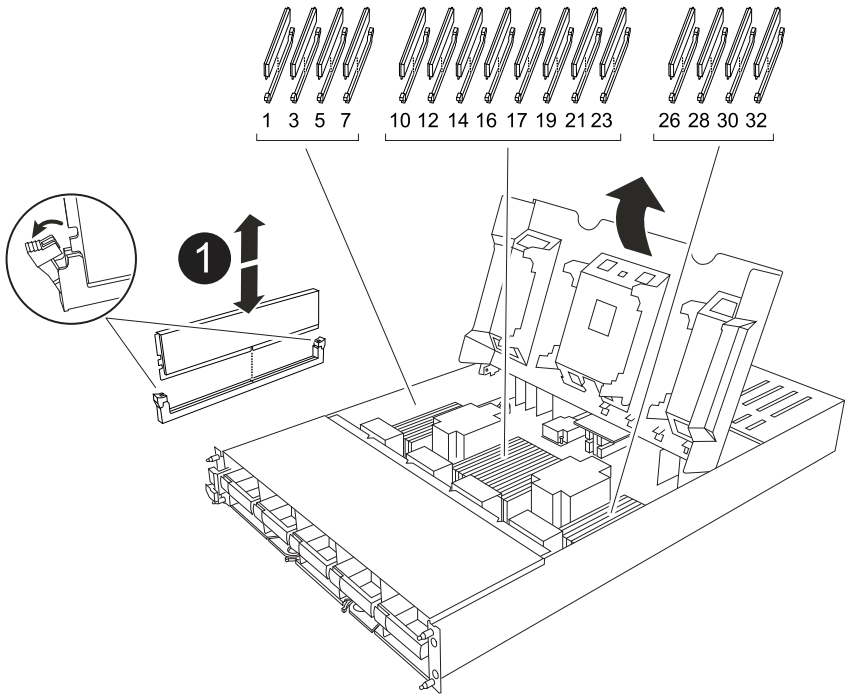

1	DIMM and DIMM ejector tabs

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the controller air duct.

Step 4: Install the controller

Reinstall the controller module and boot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch

back under the fans

> Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

The controller module begins to boot as soon as it is fully seated in the enclosure.

4. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
6. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a fan - FAS70 and FAS90

Replace a fan module in your FAS70 or FAS90 system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.

   Facing the controller module, fan modules are numbered 1 through 5, from left to right.

   > There is a single LED for each fan. It is green when the fan is functioning correctly and amber when not.

4. Press the black button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

   > The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

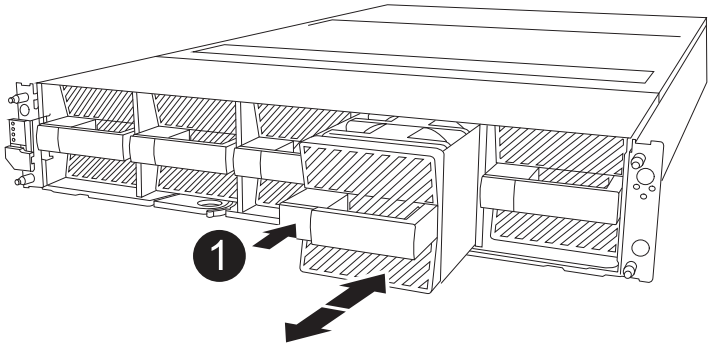

1	Black release button

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED turns off once the fan is recognized by that system.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the Flash Cache module carrier or a caching module - FAS70 and FAS90

The NVMe SSD Flash Cache module carrier in your FAS70 or FAS90 system contains one or two Flash Cache modules (caching modules) with a single SSD Flash Cache drive integrated into each caching module.

The FAS70 supports 2TB caching modules and FAS90 supports 4TB caching modules. You cannot mix caching modules of different capacity in the Flash Cache module carrier.

You can perform either of the following procedures depending on what component you need to replace: the entire Flash Cache module carrier or a caching module.

- Replace the Flash Cache module carrier
- Replace the caching module

Replace the Flash Cache module carrier

The Flash Cache module carrier is located in slot 6 and houses up to two Flash Cache modules. You cannot hot-swap the Flash Cache module carrier

Before you begin

- Ensure your storage system has the appropriate operating system for the replacement Flash Cache module carrier.
- Confirm all other components are functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired node

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the Flash Cache module carrier

Perform the following steps to replace the Flash Cache module carrier.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed Flash Cache module carrier, in slot 6, by the lit amber Attention LED on the front of the Flash Cache module carrier.

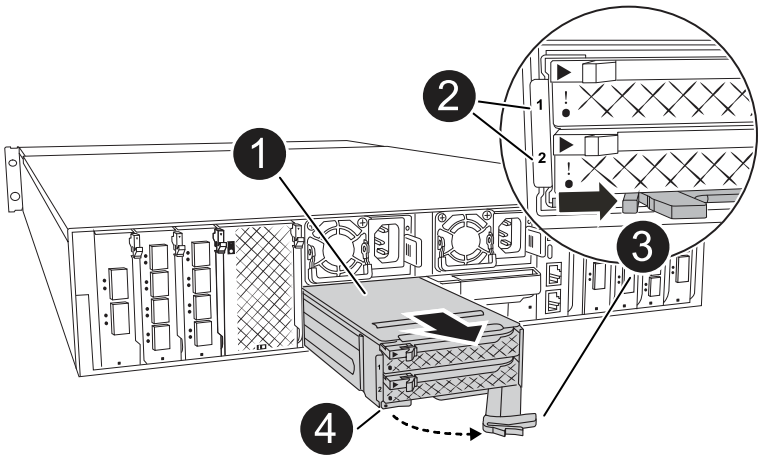

1	Flash Cache module carrier
2	Caching module slot numbers
3	Flash Cache module carrier cam handle
4	Flash Cache module carrier fault LED

3. Remove the failed Flash Cache module carrier:
   a. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

b. Pinch the blue tab at the bottom of the Flash Cache module carrier.

c. Rotate the tab away from the module.

4. Pull the Flash Cache module carrier out of the controller module and set it on an antistatic mat.
5. Move the caching modules to the replacement Flash Cache module carrier:
   a. Pinch the Terra Cotta tab at the top of the caching module and rotate the cam handle away from the caching module.
   b. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the Flash Cache module carrier.
   c. Install the caching module into the same slot in the replacement Flash Cache module carrier and rotate the cam handle to the closed position on the caching module to lock it in place.
6. Repeat these steps if there is a second caching module.
7. Install the replacement Flash Cache module carrier into the system:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
   c. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace the Flash Cache module carrier, you must reboot the controller module.

Steps

1. From the LOADER prompt, reboot the node: *bye*

   This reinitializes the I/O cards and other components and reboots the node.

2. Return the node to normal operation: *storage failover giveback -ofnode impaired_node_name*
3. If automatic giveback was disabled, reenable it: *storage failover modify -node local -auto-giveback true*

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the caching module

The Flash Cache modules (caching modules) are located in slot 6-1 or in slot 6-2 or in both slot 6-1 and slot 6-2.

You can hot-swap the individual caching modules with caching modules of the same capacity from the same vendor or from a different supported vendor.

Before you begin

- Ensure the replacement caching module has the same capacity as the failed one, from the same vendor or from a different supported vendor.
- Confirm all other components are functioning properly; if not, you must contact technical support.
- The drives in the caching modules are not field replaceable units (FRU). You must replace the entire

caching module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed caching module, in slot 6, by the lit amber Attention LED on the front of the caching module.
3. Prepare the caching module slot for replacement as follows:

   a. Record the caching module capacity, part number, and serial number on the target node: *system node run local sysconfig -av 6*

   b. In admin privilege level, prepare the target caching module slot for removal, responding `y` when prompted whether to continue: *system controller slot module remove -node node_name -slot slot_number* The following command prepares slot 6-1 on node1 for removal, and displays a message that it is safe to remove:

   ```
   ::> system controller slot module remove -node node1 -slot 6-1

   Warning: SSD module in slot 6-1 of the node node1 will be powered off
   for removal.
   Do you want to continue? (y|n): _y_
   The module has been successfully removed from service and powered
   off. It can now be safely removed.
   ```

   c. Display the slot status with the `system controller slot module show` command.

   The caching module slot status displays `powered-off` in the screen output for the caching module that needs replacing.

   See the Command man pages for your version of ONTAP for more details.

4. Remove the caching module:

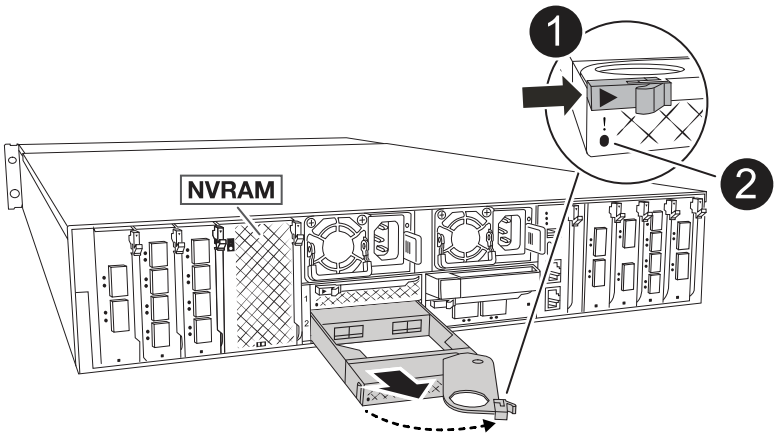

1	Caching module cam handle

2	Caching module fault LED

a. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

b. Press the terra cotta release button on the front of the caching module.

c. Rotate the cam handle as far as it will go.

d. Remove the caching module module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the Flash Cache module carrier.

   Be sure to support the caching module as you remove it from the Flash Cache module carrier.

5. Install the replacement caching module:

   a. Align the edges of the caching module with the opening in the controller module.

   b. Gently push the caching module into the bay until the cam handle engages.

   c. Rotate the cam handle until it locks into place.

   d. Rotate the cable management tray up to the closed position.

6. Bring the replacement caching module online by using the `system controller slot module insert` command as follows:

   The following command prepares slot 6-1 on node1 for power-on, and displays a message that it is powered on:

```
::> system controller slot module insert -node node1 -slot 6-1

Warning: NVMe module in slot 6-1 of the node localhost will be powered
on and initialized.
Do you want to continue? (y|n): `y`

The module has been successfully powered on, initialized and placed into
service.
```

7. Verify the slot status using the `system controller slot module show` command.

   Make sure that command output reports status for the as `powered-on` and ready for operation.

8. Verify that the replacement caching module is online and recognized, and then visually confirm that the amber attention LED is not lit: `sysconfig -av slot_number`

   If you replace the caching module with a caching module from a different vendor, the new vendor name is displayed in the command output.

9. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace NVRAM - FAS70 and FAS90

Replace the NVRAM in your FAS70 or FAS90 system when the non-volatile memory becomes faulty or requires an upgrade. The replacement process involves shutting down the impaired controller, replacing the NVRAM module or the NVRAM DIMM, reassigning the disks, and returning the failed part to NetApp.

The NVRAM module consists of the NVRAM12 hardware and field-replaceable DIMMs. You can replace a failed NVRAM module or the DIMMs inside the NVRAM module.

Before you begin

- Make sure you have the replacement part available. You must replace the failed component with a replacement component you received from NetApp.
- Make sure all other components in the storage system are functioning properly; if not, contact NetApp support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the NVRAM module or NVRAM DIMM

Replace the NVRAM module or NVRAM DIMMs using the appropriate following option.

Option 1: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 4/5 in the enclosure and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
4. Remove the impaired NVRAM module from the enclosure:
   a. Depress the locking cam button.

      The cam button moves away from the enclosure.

   b. Rotate the cam latch down as far as it will go.
   c. Remove the impaired NVRAM module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

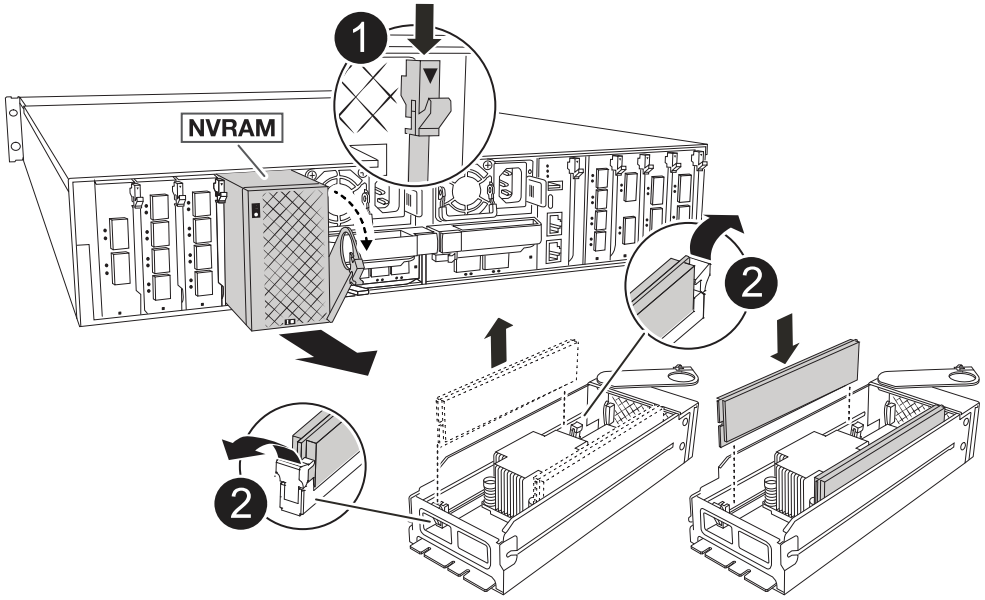

1	Cam locking button
2	DIMM locking tabs

5. Set the NVRAM module on a stable surface.
6. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
7. Install the replacement NVRAM module into the enclosure:
   a. Align the module with the edges of the enclosure opening in slot 4/5.

b. Gently slide the module into the slot all the way, and then rotate the cam latch all the way up to lock the module in place.

8. Recable the controller.
9. Rotate the cable management tray up to the closed position.

Option 2: Replace the NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, you must remove the NVRAM module, and then replace the target DIMM.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
4. Remove the target NVRAM module from the enclosure.

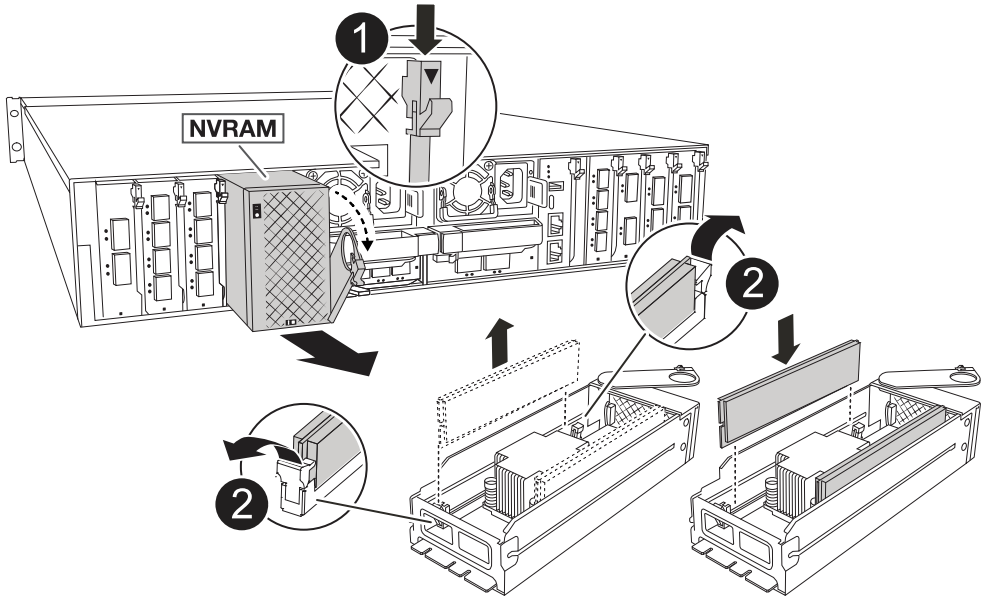

1	Cam locking button
2	DIMM locking tabs

5. Set the NVRAM module on a stable surface.
6. Locate the DIMM to be replaced inside the NVRAM module.

> Consult the FRU map label on the side of the NVRAM module to determine the locations of DIMM slots 1 and 2.

7. Remove the DIMM by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
8. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.

9. Install the NVRAM module into the enclosure:
    a. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.
10. Recable the controller.
11. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace the FRU, you must reboot the controller module by plugging the power cables back into the PSU.

Steps

1. Plug the power cables back into the PSU.

   The system will begin to reboot, typically to the LOADER prompt.

2. Enter *bye* at the LOADER prompt.
3. Return the impaired controller to normal operation by giving back its storage: *`storage failover giveback -ofnode _impaired_node_name`*.
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 4: Reassign disks

You must confirm the system ID change when you boot the controller and then verify that the change was implemented.

Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: *halt*
2. From the LOADER prompt on the controller, boot the controller and enter *y* when prompted to override the system ID due to a system ID mismatch.
3. Wait until the Waiting for giveback… message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: *storage failover show*

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1:> storage failover show
                                    Takeover
Node                Partner         Possible     State Description
------------      ------------      --------
-------------------------------------
node1               node2           false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2               node1           -            Waiting for giveback
(HA mailboxes)
```

4. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: *storage failover giveback -ofnode replacement_node_name*

      The controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter *y*.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: *storage failover show*

      The output from the `storage failover show` command should not include the System ID changed on partner message.

5. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 151759706:

```
node1:> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        151759706  151759706  -
151759706 Pool0
1.0.1  aggr0_1  node1 node1           151759706  151759706  -
151759706 Pool0
.
.
.
```

6. If the system is in a MetroCluster configuration, monitor the status of the controller: *metrocluster node show*

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the impaired system ID until the MetroCluster configuration returns to a normal state.

7. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The controller is the current owner of the disks on the disaster site.

     See Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration for more information.

8. If your system is in a MetroCluster configuration, verify that each controller is configured: *metrocluster node show - fields configuration-state*

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
--------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

9. Verify that the expected volumes are present for each controller: `vol show -node node-name`
10. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.
11. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
12. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NV battery - FAS70 and FAS90

Replace the NV battery in your FAS70 or FAS90 system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

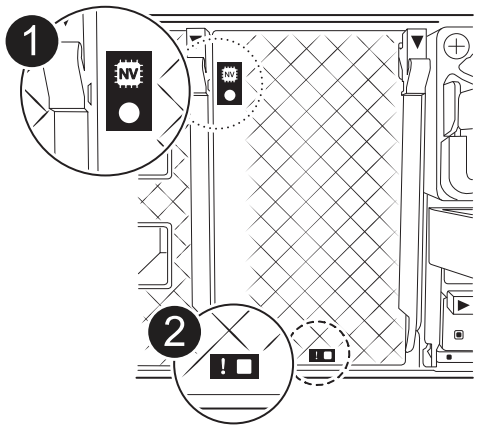

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

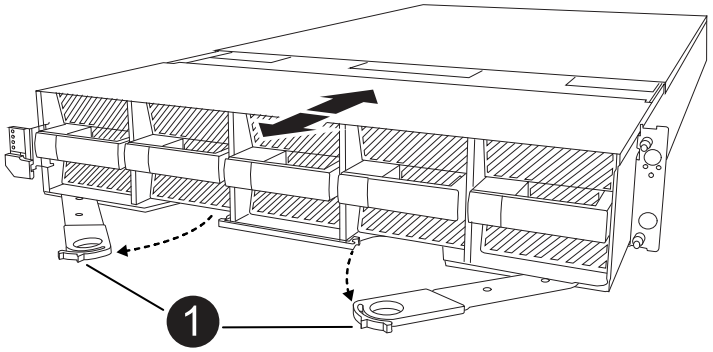

1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace the NV battery

Remove the failed NV battery from the controller module and install the replacement NV battery.

Steps

1. Open the air duct cover and locate the NV battery.

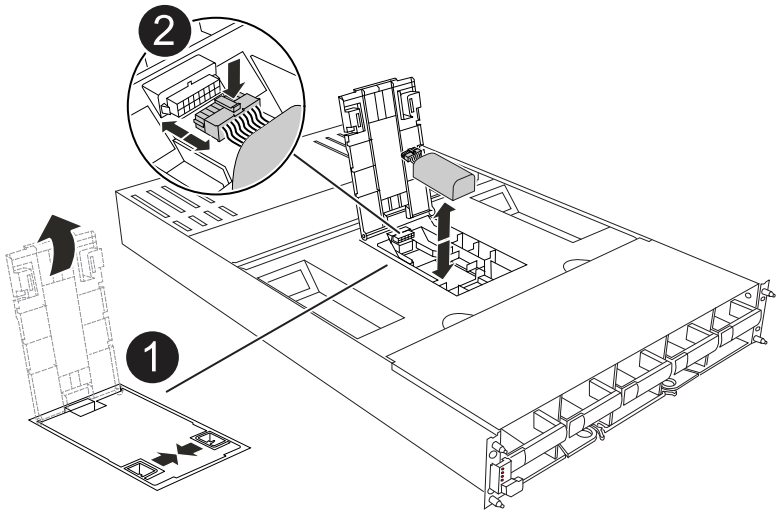

1	NV battery air duct cover
2	NV battery plug

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module, and then set it aside.

5. Remove the replacement battery from its package.
6. Install the replacement battery pack into the controller:
   a. Plug the battery plug into the riser socket and make sure that the plug locks into place.
   b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
7. Close the NV air duct cover.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

   > Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the enclosure.

4. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
6. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Overview of add and replace an I/O module - FAS70 and FAS90

The FAS70 or FAS90 system offers flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your FAS70 or FAS90 storage system with the same type of I/O module,

or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

- Add an I/O module

  Adding additional modules can improve redundancy, helping to ensure that the system remains operational even if one module fails.

- Replace an I/O module

  Replacing a failing I/O module can restore the system to its optimal operating state.

I/O slot numbering

The I/O slots on FAS70 and FAS90 controllers are numbered 1 through 11, as shown in the following illustration.

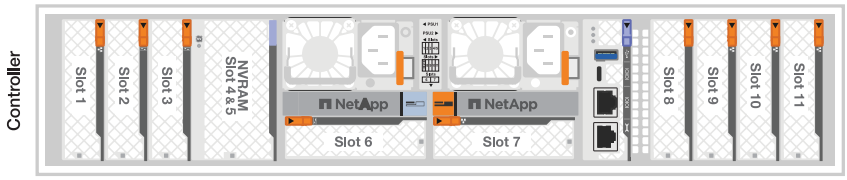

Add an I/O module - FAS70 and FAS90

Add an I/O module to yourFAS70 and FAS90 system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your FAS70 and FAS90 storage system when there are empty slots available or when all slots are fully populated.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 2: MetroCluster configuration

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster

Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- Make sure that all other components are functioning properly.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the target slot blanking module from the carrier:
   a. Depress the cam latch on the blanking module in the target slot.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
4. Install the I/O module:
   a. Align the I/O module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
5. Cable the I/O module to the designated device.

   Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

6. Rotate the cable management tray up to the closed position.
7. From the LOADER prompt, reboot the node:

   ```
   bye
   ```

   This reinitializes the I/O module and other components and reboots the node.

8. Give back the controller from the partner controller:

   ```
   storage failover giveback -ofnode target_node_name
   ```

9. Repeat these steps for controller B.
10. From the healthy node, restore automatic giveback if you disabled it:

    ```
    storage failover modify -node local -auto-giveback true
    ```

11. If AutoSupport is enabled, restore automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the target I/O module from the chassis:
   a. Depress the cam latch button.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

      Make sure that you keep track of which slot the I/O module was in.
5. Install the I/O module into the target slot in the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
6. Cable the I/O module to the designated device.
7. Repeat the remove and install steps to replace additional modules for the controller.
8. Rotate the cable management tray up to the closed position.
9. Reboot the controller from the LOADER prompt:_bye_

   This reinitializes the PCIe cards and other components and reboots the node.

   If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement
10. Give back the controller from the partner controller:

    ```
    storage failover giveback -ofnode target_node_name
    ```
11. Enable automatic giveback if it was disabled:

```
storage failover modify -node local -auto-giveback true
```

12. Do one of the following:
    - If you removed a NIC I/O module and installed a new NIC I/O module, use the following network command for each port:

      `storage port modify -node *<node name> -port *<port name> -mode network`

    - If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
13. Repeat these steps for controller B.

Replace an I/O module - FAS70 and FAS90

Replace an I/O module in your FAS70 or FAS90 system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

You can use this procedure with all versions of ONTAP supported by your storage system.

Before you begin

- You must have the replacement part available.
- Make sure all other components in the storage system are functioning properly; if not, contact technical support.

Step 1: Shut down the impaired node

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace a failed I/O module

To replace an I/O module, locate it within the enclosure and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

   This following illustration shows removing a horizontal and vertical I/O module. Typically, you will only remove one I/O module.

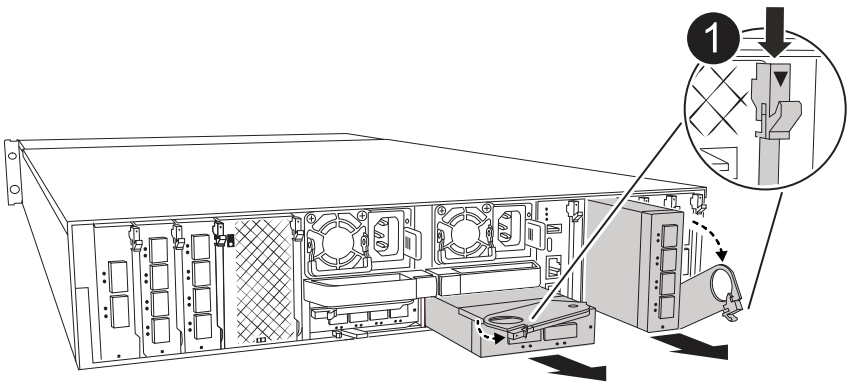

1	I/O cam latch

   Make sure that you label the cables so that you know where they came from.

4. Remove the target I/O module from the enclosure:
   a. Depress the cam button on the target module.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

      Make sure that you keep track of which slot the I/O module was in.

5. Set the I/O module aside.
6. Install the replacement I/O module into the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
7. Cable the I/O module.
8. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```

3. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```

4. If AutoSupport is enabled, restore automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - FAS70 and FAS90

Replace an AC or DC power supply unit (PSU) in your FAS70 or FAS90 system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cable, replacing the faulty PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable.

About this task

This procedure is written for replacing one PSU at a time.

Do not mix PSUs with different efficiency ratings. Always replace like for like.

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Replace an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

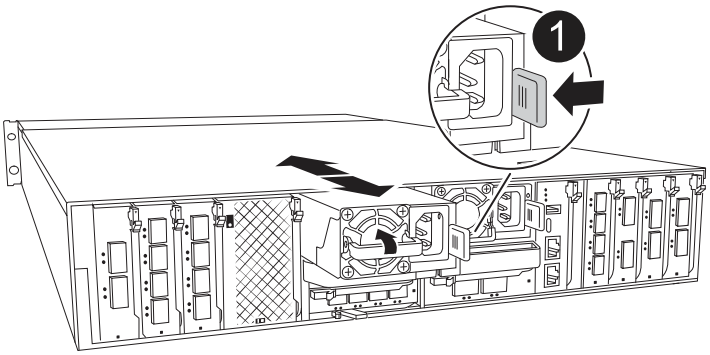

1	Terracotta PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU.
   b. Secure the power cable to the PSU using the power cable retainer.

   Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Option 2: Replace a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

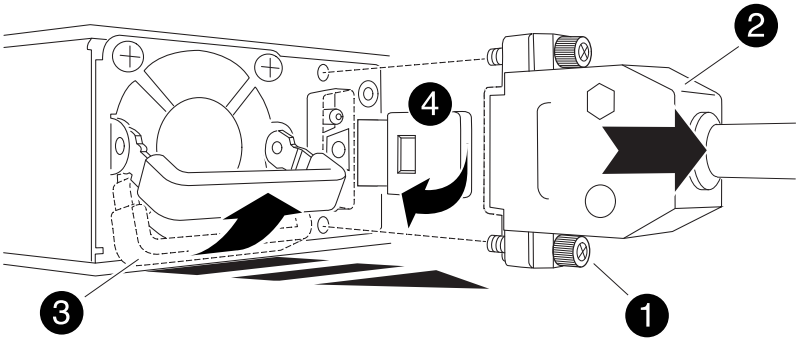

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle
4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one

way.

> To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - FAS70 and FAS90

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your FAS70 or FAS90 system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

- Understand that you can use this procedure with all versions of ONTAP supported by your system.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

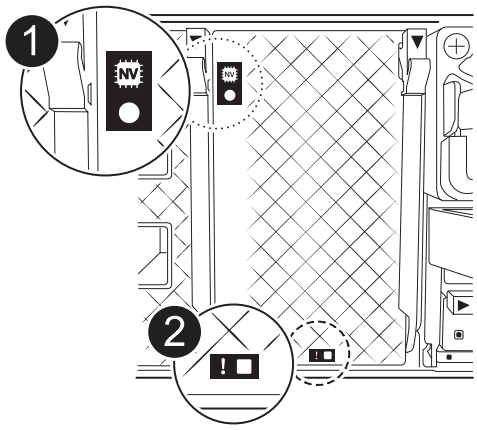

1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.

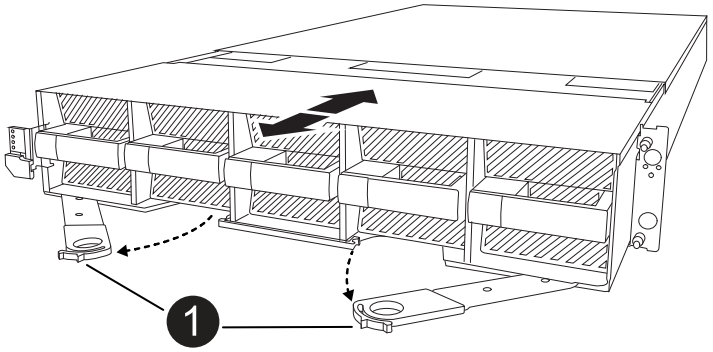

1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace the RTC battery

Remove failed RTC battery and install the replacement RTC battery.

Steps

1. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the RTC battery under the air duct.

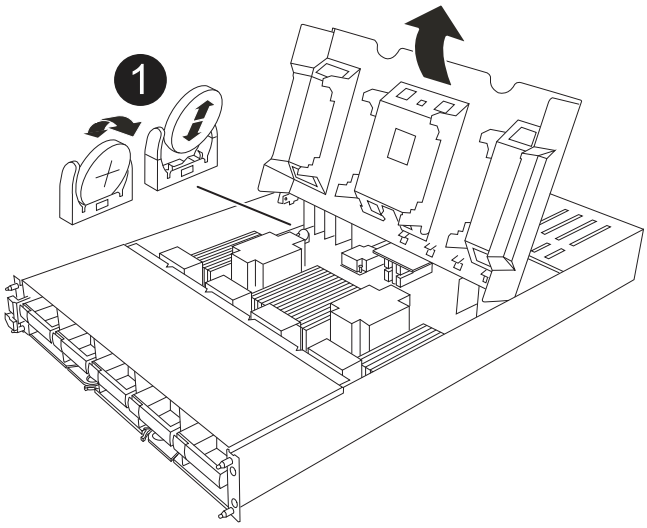

1	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

   Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the enclosure.

4. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
6. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Reset the time and date on the controller

After replacing the RTC battery, inserting controller and powering on first BIOS reset, you will see the following error messages:
`RTC date/time error. Reset date/time to default`
`RTC power failure error`
These messages are excpected and you can continue with this procedure.

Steps

1. Check the date and time on the healthy controller with the *cluster date show* command.

If your system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*

a. At the LOADER prompt on the target controller, check the time and date with the `cluster date show` command.

b. If necessary, modify the date with the `set date mm/dd/yyyy` command.

c. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.

   1. Confirm the date and time on the target controller.
   2. At the LOADER prompt, enter *bye* to reinitialize the PCIe cards and other components and let the controller reboot.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace system management module - FAS70 and FAS90

Replace the System Management module in your FAS70 or FAS90 system when it becomes defective or its firmware is corrupted. The replacement process involves shutting down the controller, replacing the failed System Management module, rebooting the controller, updating the license keys, and returning the failed part to NetApp.

The System Management module, located at the back of the controller in slot 8, contains onboard components for system management, as well as ports for external management. The target controller must be shut down to replace an impaired System Management module or replace the boot media.

The System Management module has the following onboard components:

- Boot media, allowing boot media replacement without removing the controller module.
- BMC
- Management switch

The System Management module also contains the following ports for external management:

- RJ45 Serial
- USB Serial (Type-C)
- USB Type-A (Boot recovery)
- e0M RJ45 Ethernet

Before you begin

- Make sure all other system components are working properly.
- Make sure that the partner controller is able to take over the impaired controller.
- Make sure you replace the failed component with a replacement component you received from NetApp.

About this task

This procedure uses the following terminology:

- The impaired controller is the controller on which you are performing maintenance.
- The healthy controller is the HA partner of the impaired controller.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the impaired System Management module

Replace the impaired system management module.

Steps

1. Remove the System Management module:

> Make sure NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged. If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

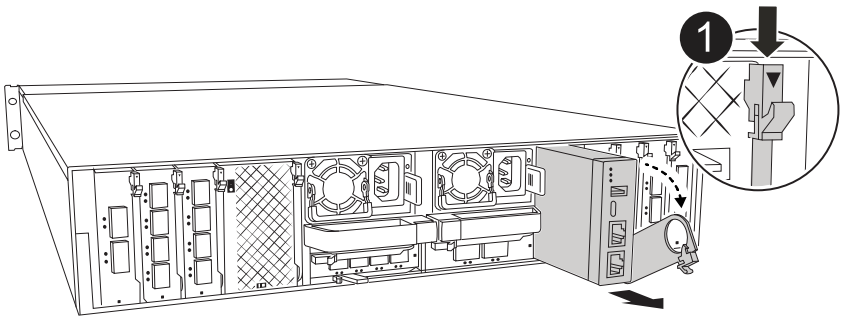

1	System Management module cam latch

   a. If you are not already grounded, properly ground yourself.
   b. Unplug the power supply cables from the PSUs.
2. Remove the System Manage module
   a. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Disconnect the power cords from the PSU for the impaired controller.
   c. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   d. Depress the cam button on the System Management module.
   e. Rotate the cam lever down as far as it will go.

f. Loop your finger into the hole on the cam lever and pull the module straight out of the system.

g. Place the System Management module on an anti-static mat, so that the boot media is accessible.

3. Move the boot media to the replacement System Management module:

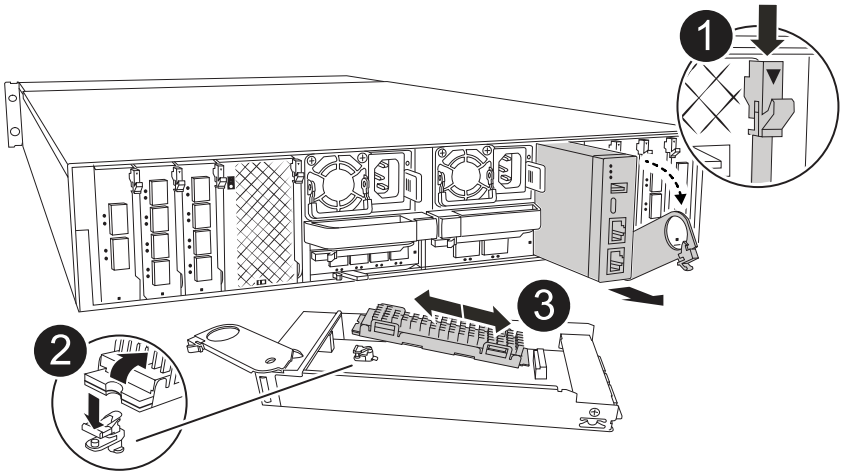

1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue boot media locking button in the impaired System Management module.

b. Rotate the boot media up and slide it out of the socket.

4. Install the boot media in the replacement System Management module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down until it touches the locking button.

c. Depress the blue locking and rotate the boot media all the way down and release the blue locking button.

5. Install the replacement System Management module into the enclosure:

a. Align the edges of the replacement System Management module with the system opening and gently push it into the controller module.

b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

6. Rotate the cable management arm up to the closed position.

7. Recable the System Management module.

Step 3: Reboot the controller module

Reboot the controller module.

Steps

1. Plug the power cables back into the PSU.

   The system will begin to reboot, typically to the LOADER prompt.

2. Enter *bye* at the LOADER prompt.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 4: Install licenses and register serial number

You must install new licenses for the node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the node. However, if the node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the for the node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

4. Register the system serial number with NetApp Support.
    - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
    - If AutoSupport is not enabled, call NetApp Support to register the serial number.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

FAS2700 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick guide - FAS2700

This page gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Access the *Installation and Setup Instructions* PDF poster:

AFF A220/FAS2700 Systems Installation and Setup Instructions

Video steps - FAS2700

The following video shows how to install and cable your new system.

[] | *https://img.youtube.com/vi/FUtG1Je5D1g?/maxresdefault.jpg*

Detailed guide - FAS2700

This page gives detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your FAS2700 system, you need to create an account on the NetApp Support Site, register your system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser
- A laptop or console with an RJ-45 connection and access to a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register your system.

      NetApp Product Registration

4. Download and install Config Advisor on your laptop.

   NetApp Downloads: Config Advisor

5. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
10 GbE cable (order dependent)	X6566B-05-R6 (112-00297), 0.5m X6566B-2-R6 (112-00299), 2m		Cluster interconnect network

Type of cable…	Part number and length	Connector type	For…
10 GbE cable (order dependent)	Part number X6566B-2-R6 (112-00299), 2m or X6566B-3-R6 (112-00300), 3m X6566B-5-R6 (112-00301), 5m		Data
Optical network cables (order dependent)	X6553-R6 (112-00188), 2m X6536-R6 (112-00090), 5m X6554-R6(112-00189), 15m		FC host network
Cat 6, RJ-45 (order dependent)	Part numbers X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network and Ethernet data
Storage (order dependent)	Part number X66030A (112-00435), 0.5m X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		Storage
Micro-USB console cable	Not applicable		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	Not applicable		Powering up the system

6. Download and complete the *Cluster configuration worksheet*.

   Cluster Configuration Worksheet

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

3. Attach cable management devices (as shown).

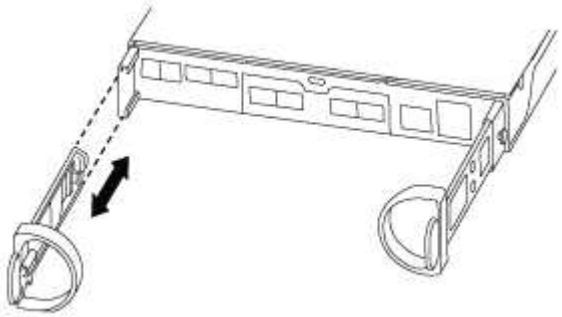

4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Cable a two-node switchless cluster, unified network configuration

Management network, UTA2 data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

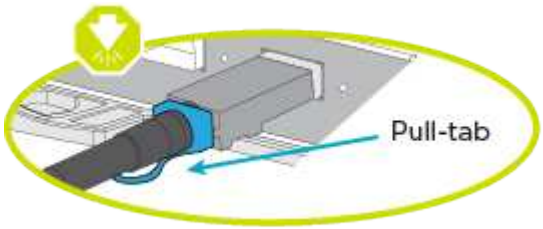

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and to the switches:

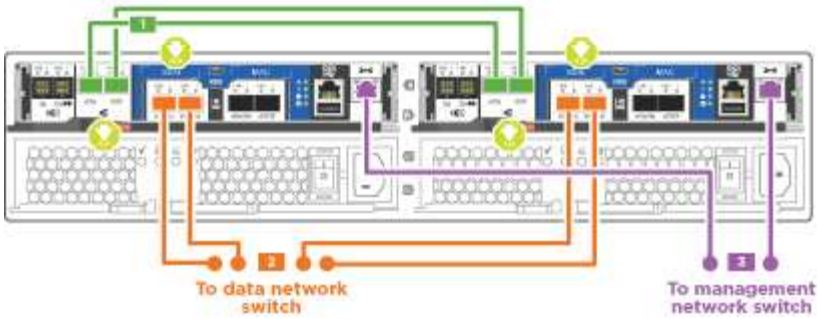

Step	Perform on each controller
1	Cable the cluster interconnect ports to each other with the cluster interconnect cable: • e0a to e0a • e0b to e0b  Cluster interconnect cables 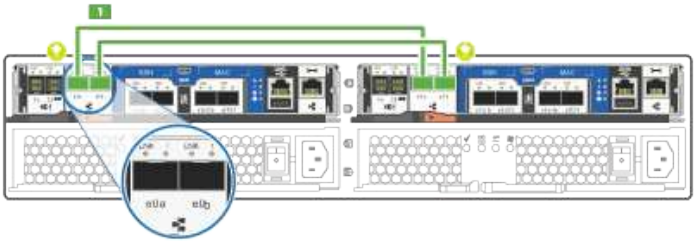

Step	Perform on each controller
2	Use one of the following cable types to cable the UTA2 data ports to your host network: An FC host • 0c and 0d • **or** 0e and 0f A 10GbE • e0c and e0d • **or** e0e and e0f You can connect one port pair as CNA and one port pair as FC, or you can connect both port pairs as CNA or both port pairs as FC. 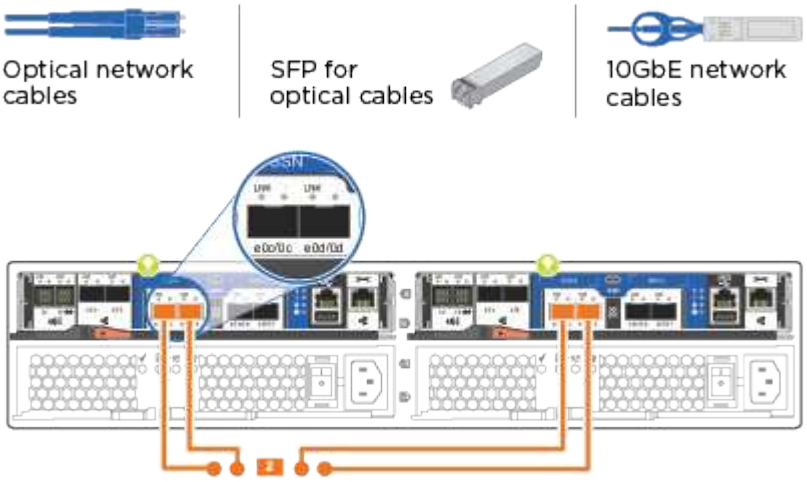
3	Cable the e0M ports to the management network switches with the RJ45 cables: 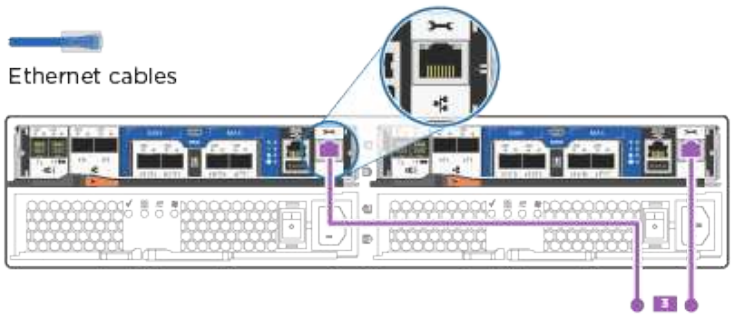
	DO NOT plug in the power cords at this point.

2. To cable your storage, see Step 4: Cable controllers to drive shelves

Option 2: Cable a switched cluster, unified network configuration

Management network, UTA2 data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled to the cluster interconnect switches.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

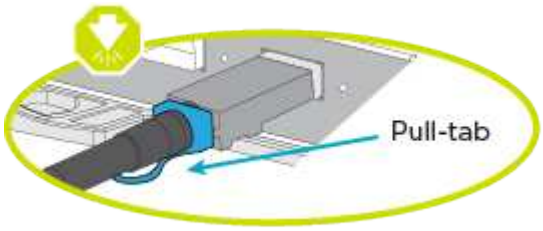

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and the switches:

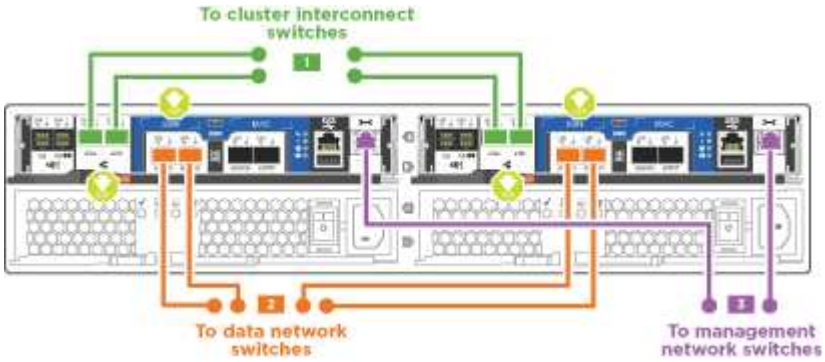

Step	Perform on each controller module
1	Cable e0a and e0b to the cluster interconnect switches with the cluster interconnect cable: 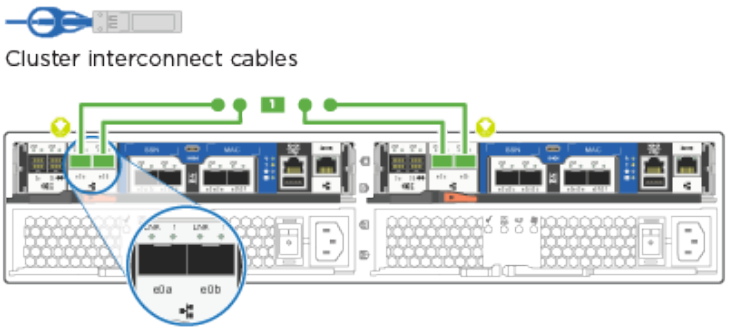
2	Use one of the following cable types to cable the UTA2 data ports to your host network: An FC host • 0c and 0d • **or** 0e and 0f A 10GbE • e0c and e0d • **or** e0e and e0f You can connect one port pair as CNA and one port pair as FC, or you can connect both port pairs as CNA or both port pairs as FC. 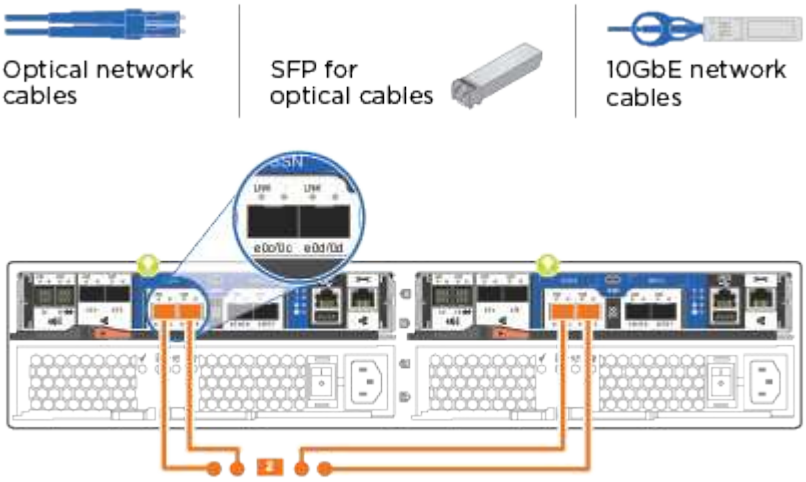

Step	Perform on each controller module
3	Cable the e0M ports to the management network switches with the RJ45 cables: 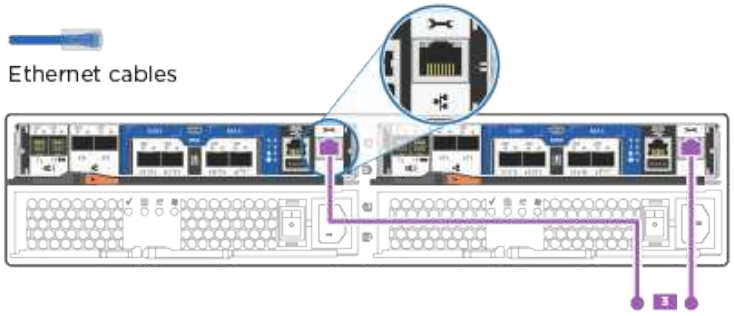
	DO NOT plug in the power cords at this point.

2. To cable your storage, see Step 4: Cable controllers to drive shelves

Option 3: Cable a two-node switchless cluster, Ethernet network configuration

Management network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

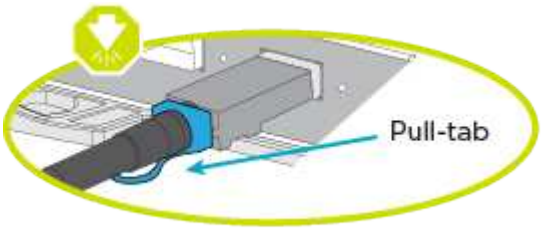

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and to the switches:

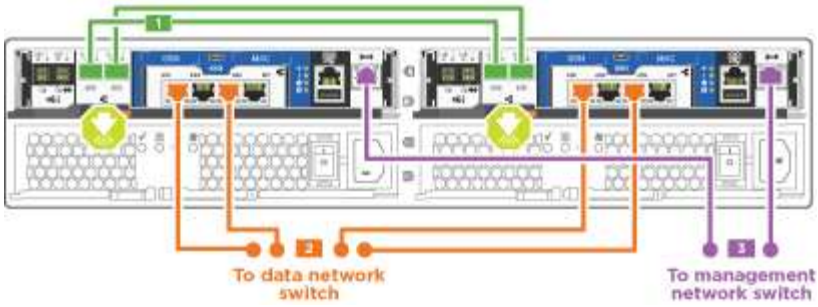

Step	Perform on each controller
1	Cable the cluster interconnect ports to each other with the cluster interconnect cable: • e0a to e0a • e0b to e0b 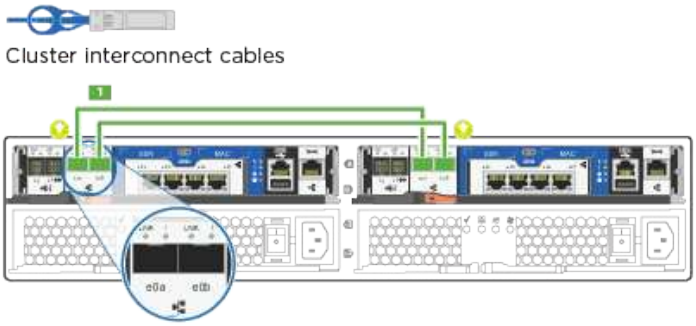
2	Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network: 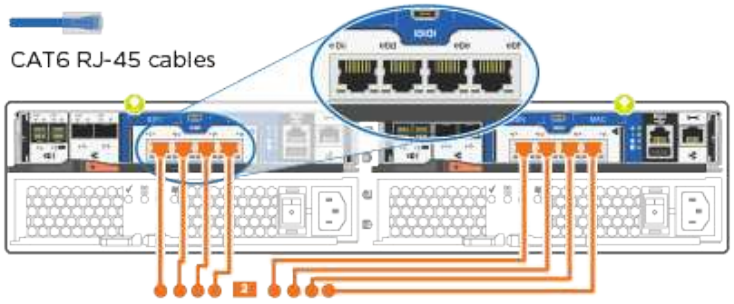

Step	Perform on each controller
3	Cable the e0M ports to the management network switches with the RJ45 cables: 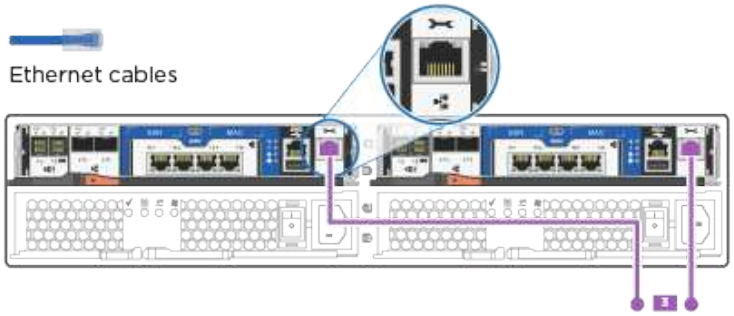
!	DO NOT plug in the power cords at this point.

2. To cable your storage, see Step 4: Cable controllers to drive shelves

Option 4: Cable a switched cluster, Ethernet network configuration

Management network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled to the cluster interconnect switches.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

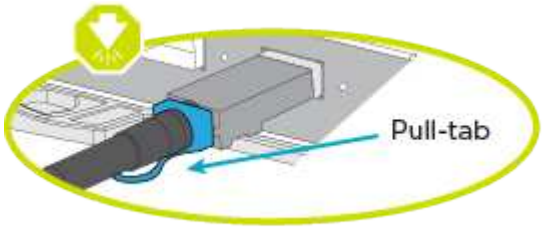

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and the switches:

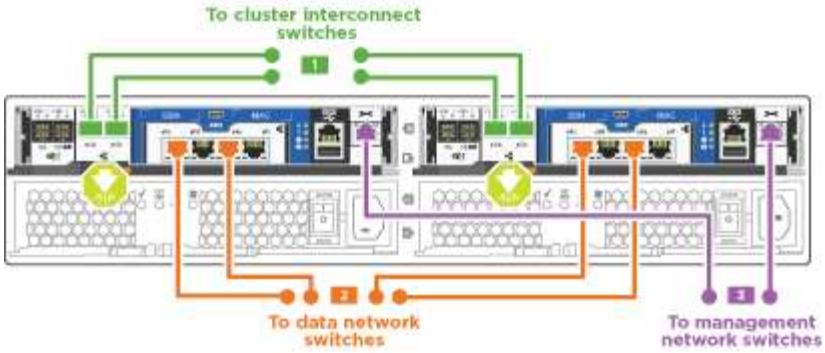

Step	Perform on each controller module
1	Cable e0a and e0b to the cluster interconnect switches with the cluster interconnect cable: 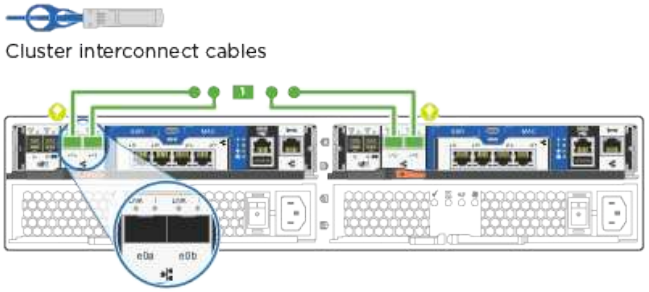
2	Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network: 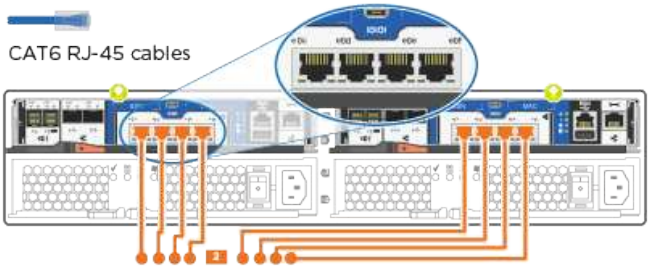

Step	Perform on each controller module
3	Cable the e0M ports to the management network switches with the RJ45 cables: 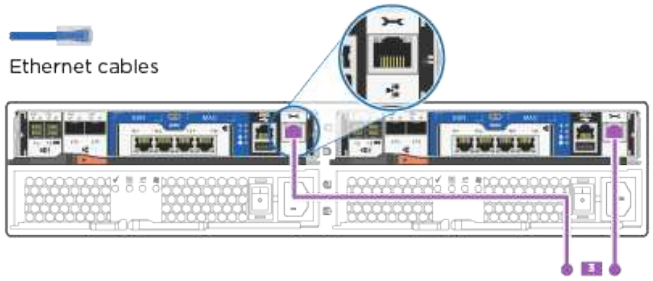
	DO NOT plug in the power cords at this point.

2. To cable your storage, see Step 4: Cable controllers to drive shelves

Step 4: Cable controllers to drive shelves

You must cable the controllers to your shelves using the onboard storage ports. NetApp recommends MP-HA cabling for systems with external storage. If you have a SAS tape drive, you can use single-path cabling. If you have no external shelves, MP-HA cabling to internal drives is optional (not shown) if the SAS cables are ordered with the system.

Before you begin

- You must cable the shelf-to-shelf connections, and then cable both controllers to the drive shelves.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

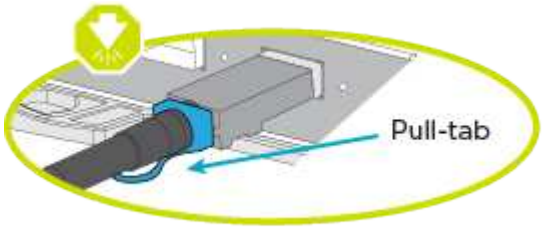

Steps

1. Cable the HA pair with external drive shelves:

The example uses DS224C. Cabling is similar with other supported drive shelves.

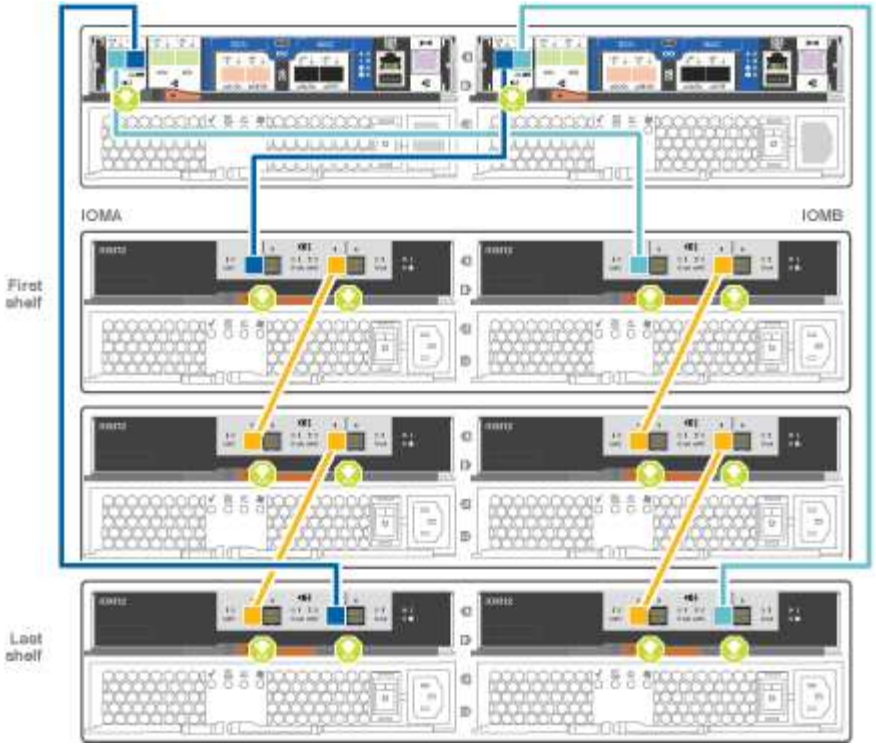

Step	Perform on each controller
1	Cable the shelf-to-shelf ports. • Port 3 on IOM A to port 1 on the IOM A on the shelf directly below. • Port 3 on IOM B to port 1 on the IOM B on the shelf directly below.  mini-SAS HD to mini-SAS HD cables
2	Connect each node to IOM A in the stack. • Controller 1 port 0b to IOM A port 3 on last drive shelf in the stack. • Controller 2 port 0a to IOM A port 1 on the first drive shelf in the stack.  mini-SAS HD to mini-SAS HD cables
3	Connect each node to IOM B in the stack • Controller 1 port 0a to IOM B port 1 on first drive shelf in the stack. • Controller 2 port 0b to IOM B port 3 on the last drive shelf in the stack.  mini-SAS HD to mini-SAS HD cables

For more SAS cabling information and worksheets, see SAS cabling rules, worksheets, and examples overview - shelves with IOM12 modules

1. To complete setting up your system, see Step 5: Complete system setup and configuration

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Complete system setup if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Use the following animation to set one or more drive shelf IDs

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Turn on the power switches to both nodes.

   Initial booting may take up to eight minutes.

4. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

5. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

6. Select an ONTAP icon listed to discover:

   drw_autodiscovery_controler_select_ieops-1849.svg[Select an ONTAP icon]

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

   System Manager opens.

7. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

ONTAP Configuration Guide

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your system.

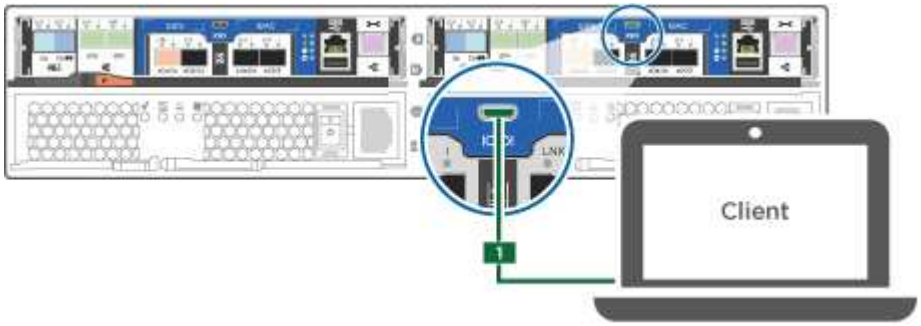

   c. Connect the laptop or console to the switch on the management subnet.

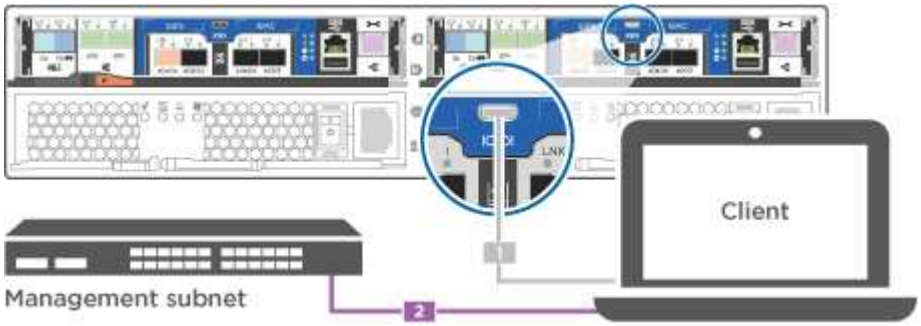

   d. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Use the following animation to set one or more drive shelf IDs:

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
4. Turn on the power switches to both nodes.

Initial booting may take up to eight minutes.

5. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

6. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.
      The format for the address is https://x.x.x.x.
   b. Configure the system using the data you collected in the *NetApp ONTAP Configuration guide*.
      ONTAP Configuration Guide
7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain FAS2700 hardware

For the FAS2700 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Caching module

You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

NVMEM battery

A battery is included with the controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real-time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF A220 and FAS2700

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:

    - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
    - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct node:
    - The *impaired* node is the node on which you are performing maintenance.
    - The *healthy node* is the HA partner of the impaired node.

Check encryption key support and status - AFF A220 and FAS2700

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
    - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
    - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the impaired controller - AFF A220 and FAS2700

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller displays…	Then…
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - FAS2700

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

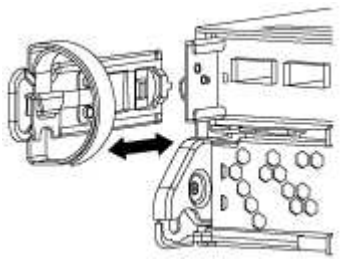

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

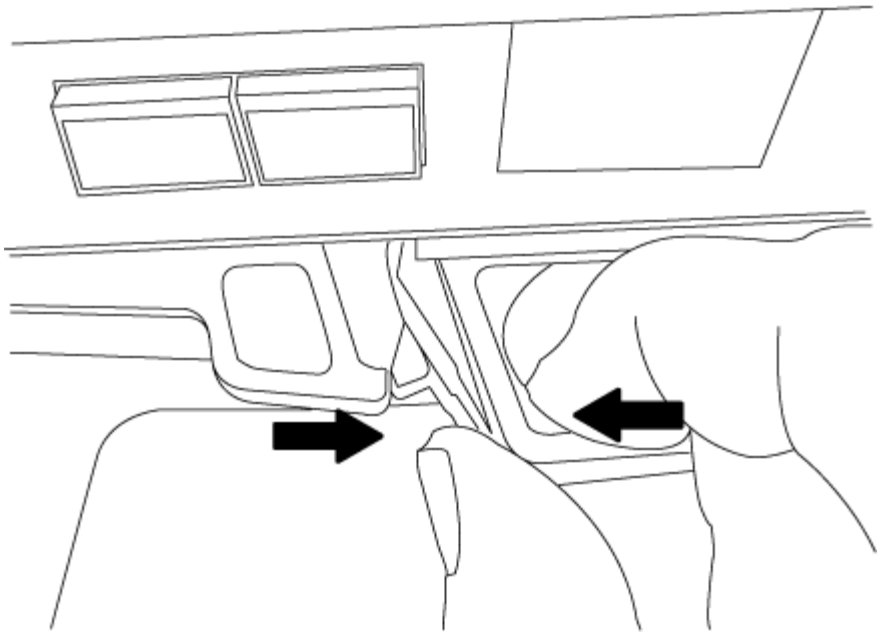

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

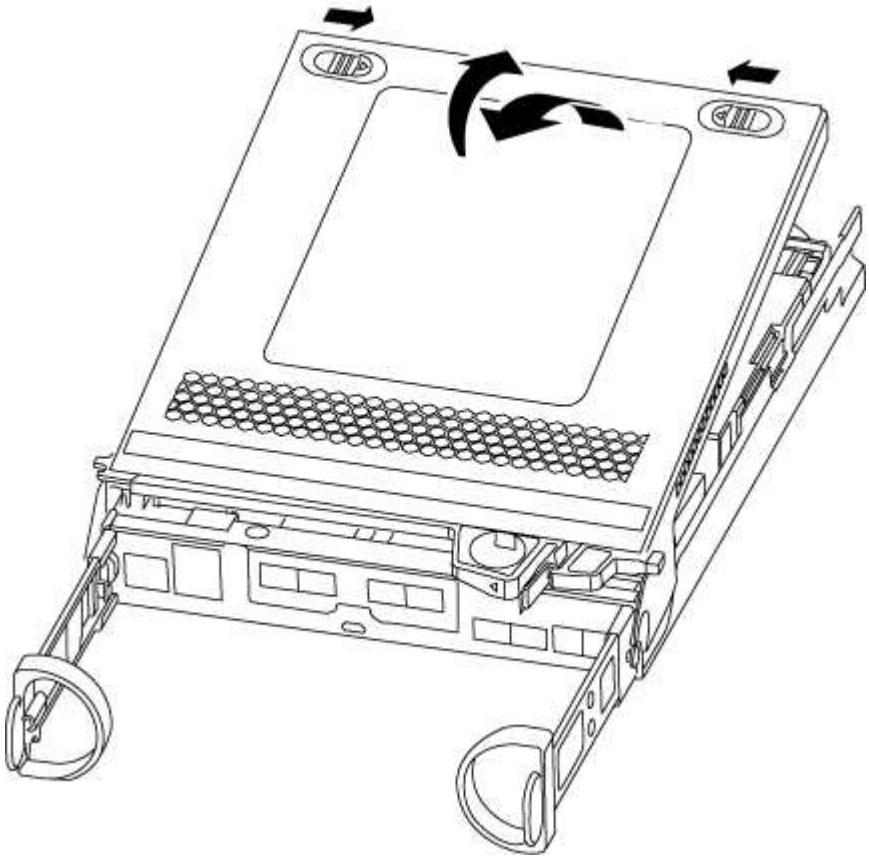

Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the boot media using the following illustration or the FRU map on the controller module:

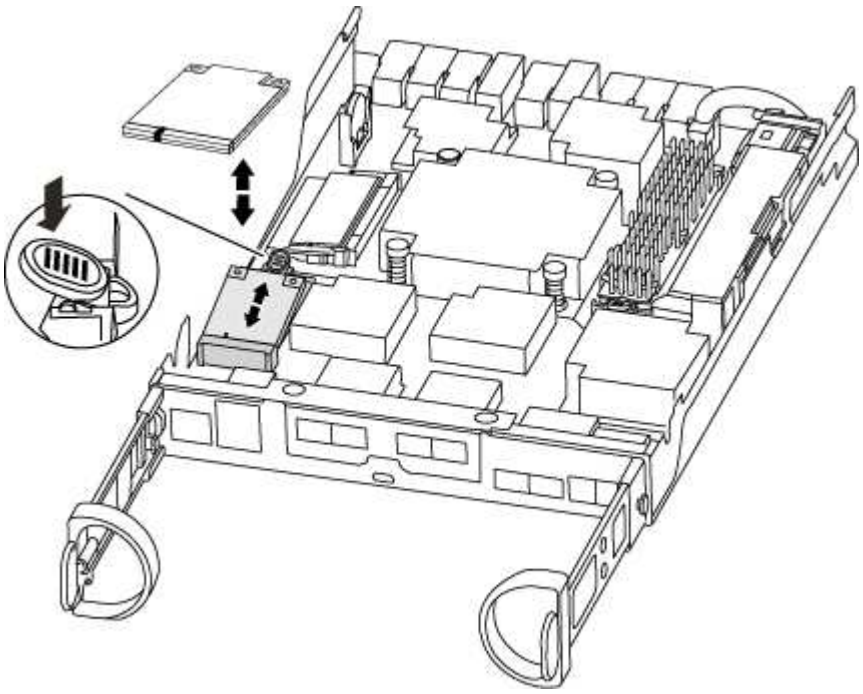

3. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

4. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
5. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

6. Push the boot media down to engage the locking button on the boot media housing.
7. Close the controller module cover.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

   The controller begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

6. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

   The system begins to boot and stops at the LOADER prompt.

7. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

     The target port you configure is the target port you use to communicate with the impaired controller from the healthy controller during var file system restore with a network connection. You can also use the e0M port in this command.

   - If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
     - filer_addr is the IP address of the storage system.
     - netmask is the network mask of the management network that is connected to the HA partner.
     - gateway is the gateway for the network.
     - dns_addr is the IP address of a name server on your network.
     - dns_domain is the Domain Name System (DNS) domain name.

       If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

Boot the recovery image - FAS2700

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:

- If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed
- If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.
8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore encryption - AFF A220 and FAS2700

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.

- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF A220 and FAS2700

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the caching module - FAS2700

You must replace the caching module in the controller module when your system registers a single AutoSupport (ASUP) message that the module has gone offline; failure to do so results in performance degradation.

- You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller.

Synchronize a node with the cluster

You might want to erase the contents of your caching module before replacing it.

Steps

1. Although data on the caching module is encrypted, you might want to erase any data from the impaired caching module and verify that the caching module has no data:
   a. Erase the data on the caching module: `system controller flash-cache secure-erase run -node node name localhost -device-id device_number`

      Run the `system controller flash-cache show` command if you don't know the Flash Cache device ID.

   b. Verify that the data has been erased from the caching module: `system controller flash-cache secure-erase show`
2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller: • For an HA pair, take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`. • For a stand-alone system: `system node halt impaired_node_name`

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

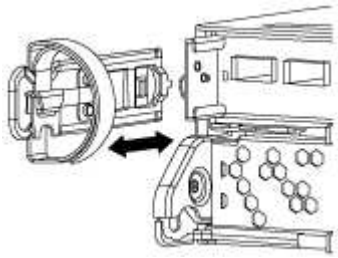

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

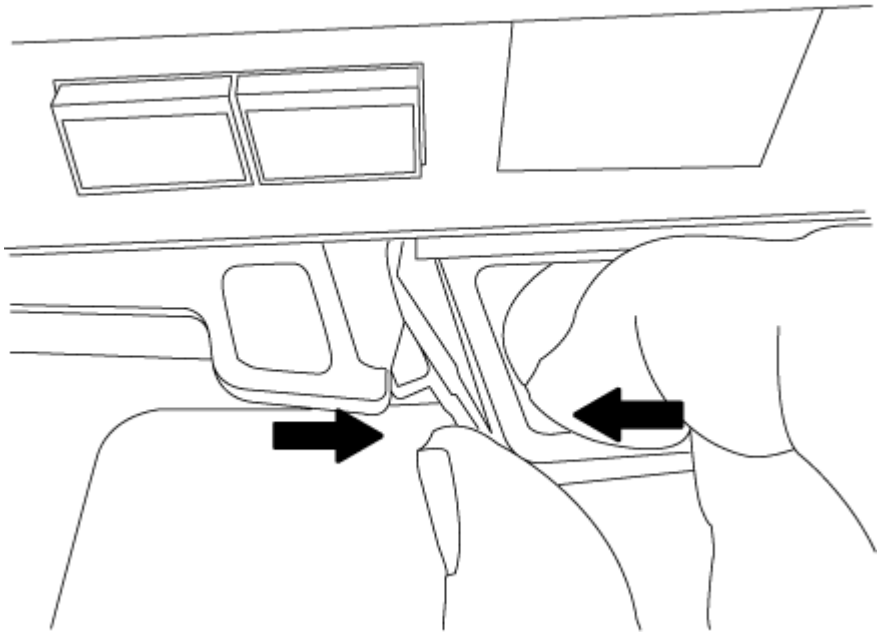

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

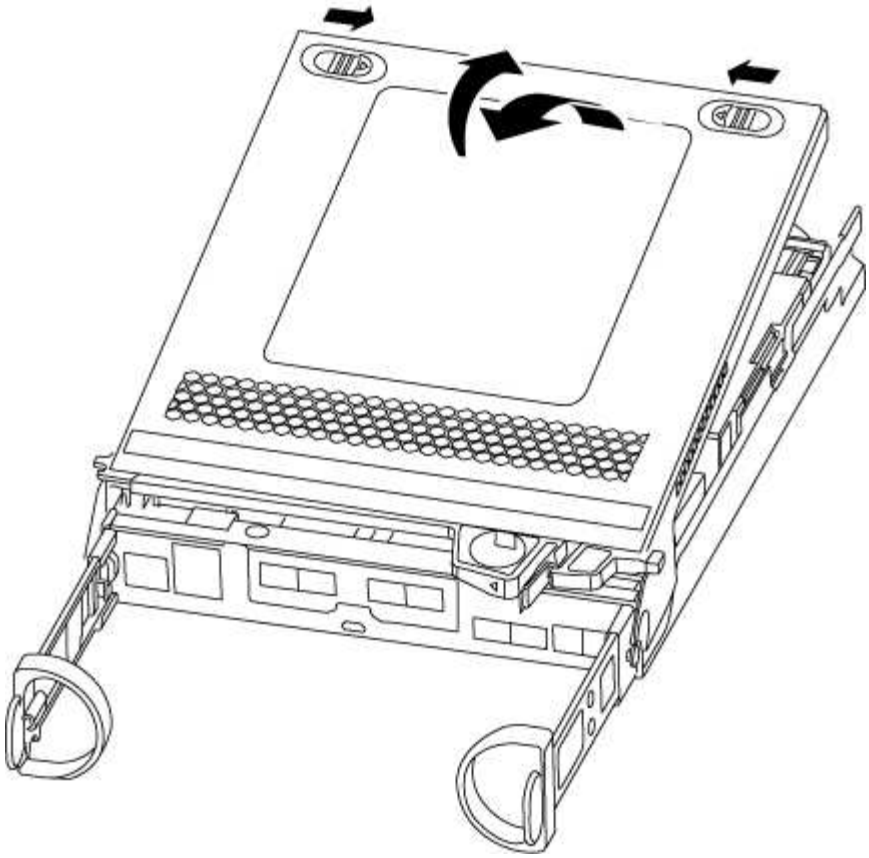

Step 3: Replace a caching module

To replace a caching module referred to as the M.2 PCIe card on the label on your controller, locate the slot inside the controller and follow the specific sequence of steps.

Your storage system must meet certain criteria depending on your situation:

- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- All other components in the storage system must be functioning properly; if not, you must contact technical support.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the caching module at the rear of the controller module and remove it.
    a. Press the release tab.
    b. Remove the heatsink.

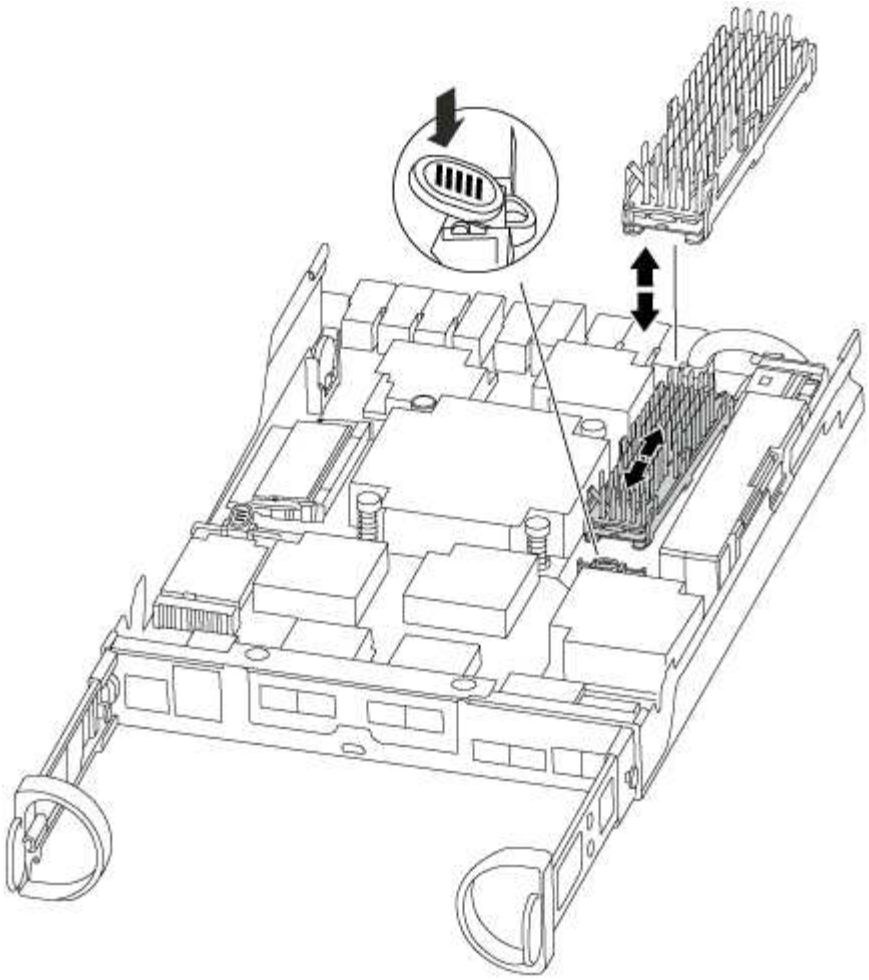

3. Gently pull the caching module straight out of the housing.
4. Align the edges of the caching module with the socket in the housing, and then gently push it into the socket.
5. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

6. Reseat and push the heatsink down to engage the locking button on the caching module housing.
7. Close the controller module cover, as needed.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - FAS2700

To replace the chassis, you must move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving all drives and controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-controller cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - FAS2700

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.

3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Controller is in a MetroCluster configuration

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message`

`MAINT=number_of_hours_downh`

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Move and replace hardware - AFF A220 and FAS2700

Move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Move a power supply

Moving out a power supply when replacing a chassis involves turning off, disconnecting, and removing the power supply from the old chassis and installing and connecting it on the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.
4. Use the cam handle to slide the power supply out of the system.

   

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.

6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   > Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

   > Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

Step 2: Remove the controller module

Remove the controller module or modules from the old chassis.

1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

2. Remove and set aside the cable management devices from the left and right sides of the controller module.

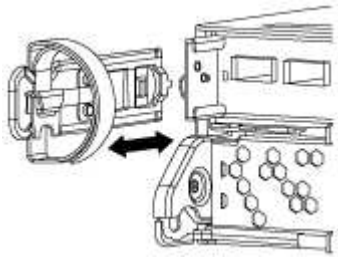

3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

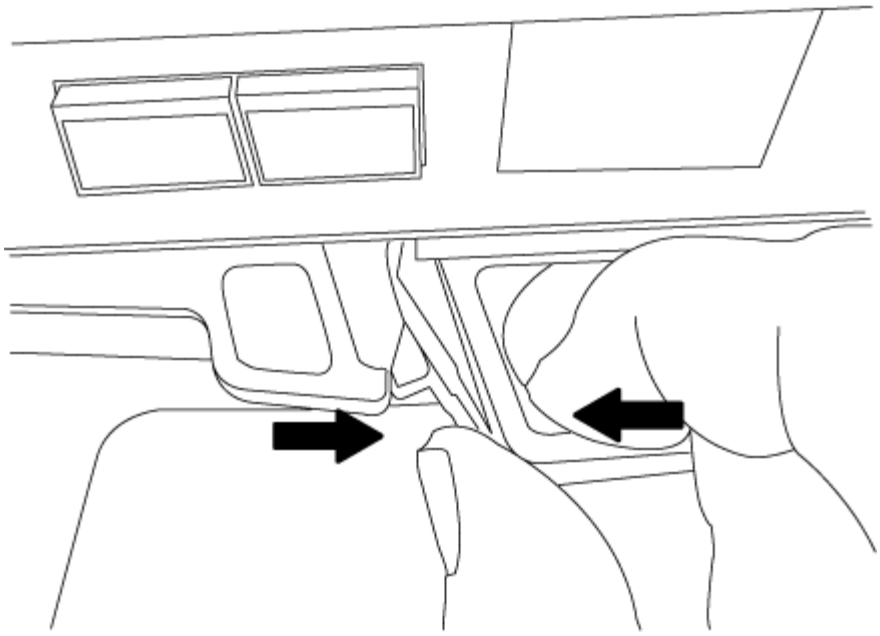

4. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 3: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.

5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It click when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 4: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller

After you install the controller module and any other components into the new chassis, boot it the system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Repeat the preceding steps if there is a second controller to install in the new chassis.
4. Complete the installation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Repeat the preceding steps for the second controller module in the new chassis.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reinstall the blanking panel and then go to the next step.

5. Connect the power supplies to different power sources, and then turn them on.
6. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

      If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

   b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - FAS2700

You must verify the HA state of the chassis, switch back aggregates, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`

- `mcc-2n`
- `mccip`
- `non-ha`

b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.
4. The next step depends on your system configuration.
5. Reboot the system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - FAS2700

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- This procedure includes steps for automatically or manually reassigning drives to the *replacement* controller, depending on your system’s configuration.

  You should perform the drive reassignment as directed in the procedure.

- You must replace the failed component with a replacement FRU component you received from your

provider.

- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - FAS2700

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - FAS2700

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove controller module

To replace the controller module, you must first remove the old controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

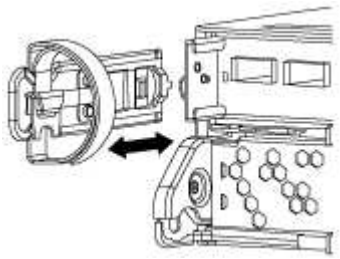

4. If you left the SFP modules in the system after removing the cables, move them to the new controller module.
5. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

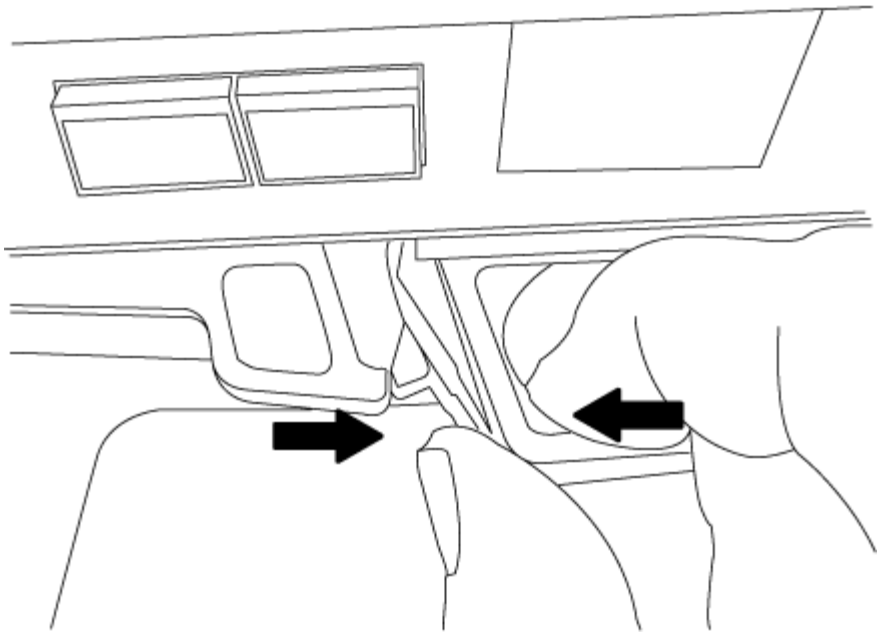

6. Turn the controller module over and place it on a flat, stable surface.
7. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

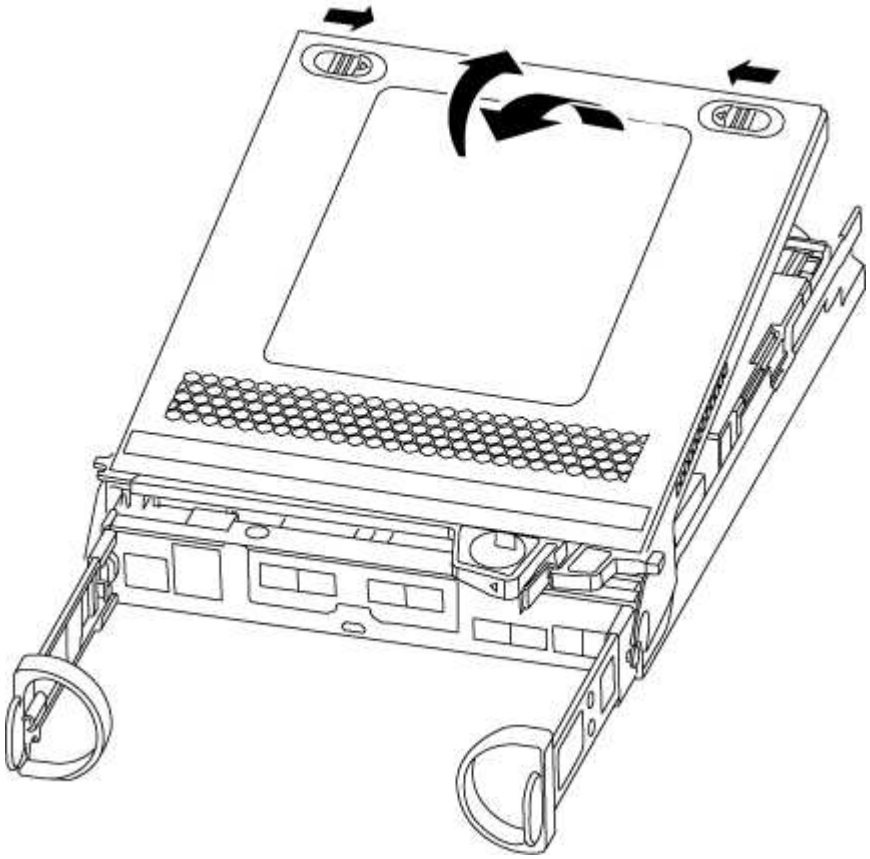

Step 2: Move the NVMEM battery

To move the NVMEM battery from the old controller module to the new controller module, you must perform a specific sequence of steps.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Locate the NVMEM battery in the controller module.

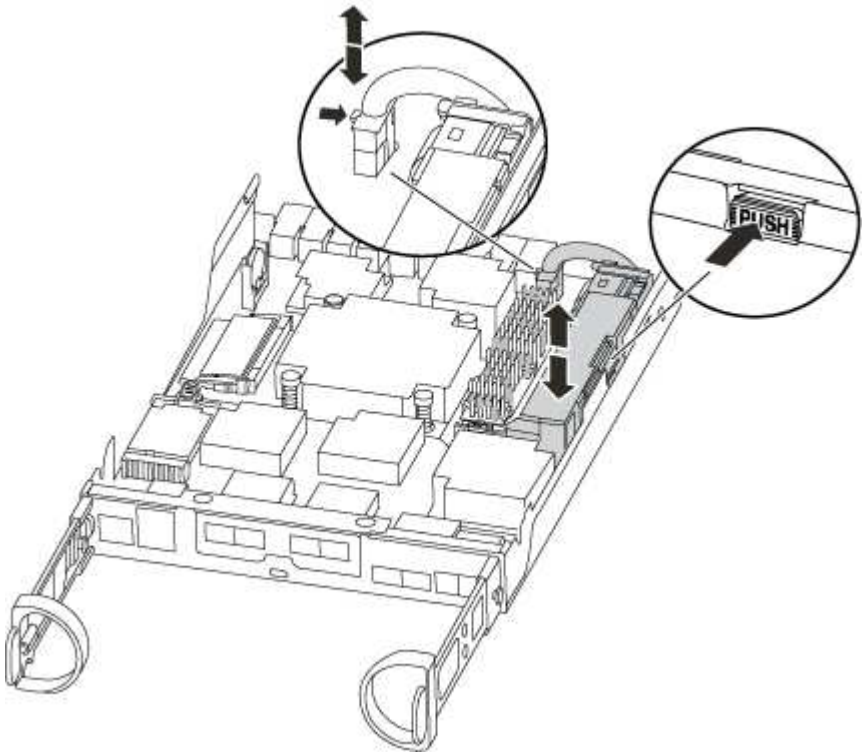

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Loop the battery cable around the cable channel on the side of the battery holder.

7. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
8. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.

Step 3: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller module and insert it in the new controller module.

1. Locate the boot media using the following illustration or the FRU map on the controller module:

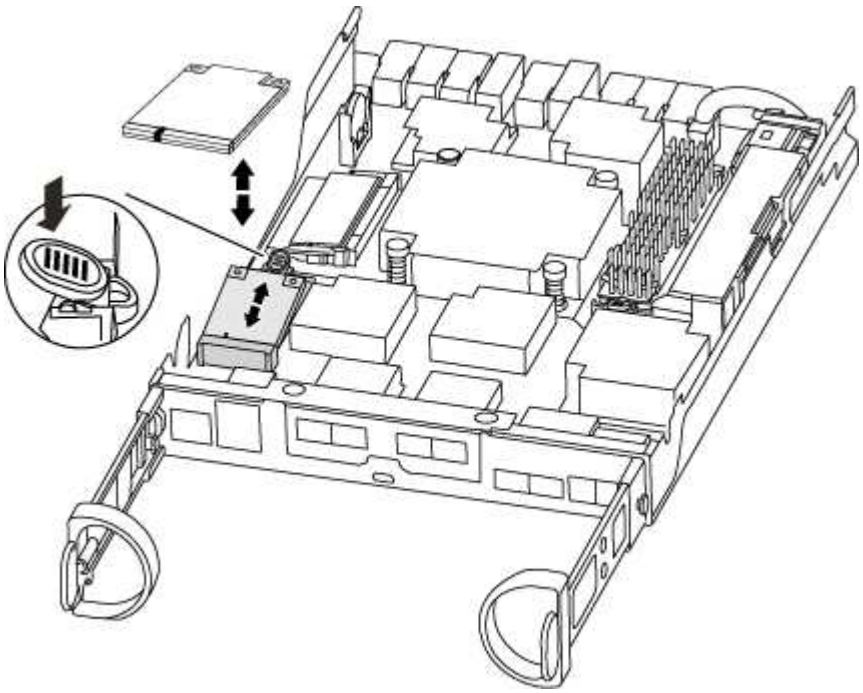

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 4: Move the DIMMs

To move the DIMMs, you must follow the directions to locate and move them from the old controller module into the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired

controller module to the corresponding slots in the replacement controller module.

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   The number and placement of system DIMMs depends on the model of your system.

   The following illustration shows the location of system DIMMs:

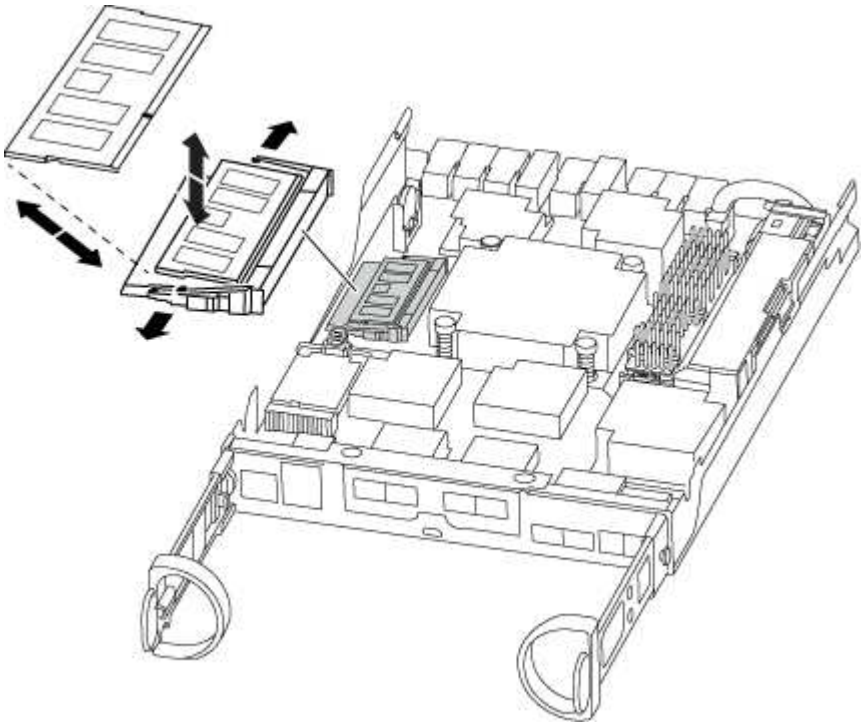

4. Repeat these steps to remove additional DIMMs as needed.
5. Verify that the NVMEM battery is not plugged into the new controller module.
6. Locate the slot where you are installing the DIMM.
7. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Repeat these steps for the remaining DIMMs.
9. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

Make sure that the plug locks down onto the controller module.

Step 5: Move a caching module, if present

If your AFF A220 or FAS2700 system has a caching module, you need to move the caching module from the old controller module to the replacement controller module. The caching module is referred to as the “M.2 PCIe card” on the controller module label.

You must have the new controller module ready so that you can move the caching module directly from the old controller module to the corresponding slot in the new one. All other components in the storage system must be functioning properly; if not, you must contact technical support.

1. Locate the caching module at the rear of the controller module and remove it.

   a. Press the release tab.

   b. Remove the heatsink.

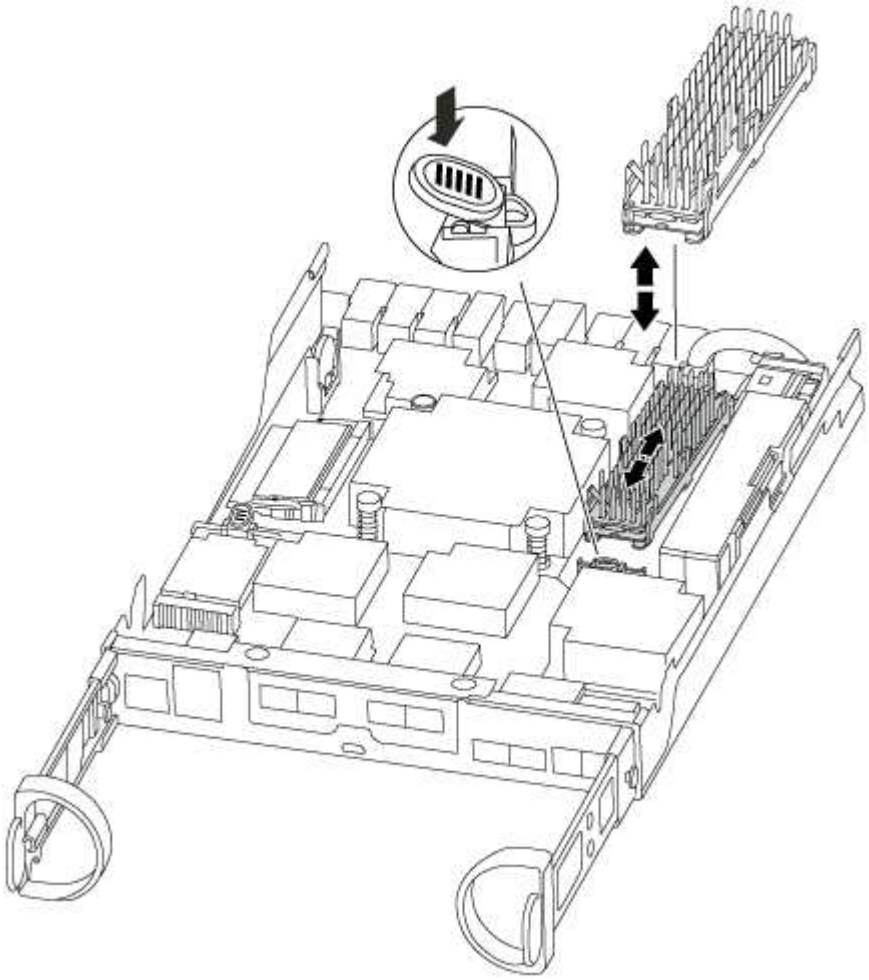

2. Gently pull the caching module straight out of the housing.
3. Move the caching module to the new controller module, and then align the edges of the caching module with the socket housing and gently push it into the socket.
4. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

5. Reseat and push the heatsink down to engage the locking button on the caching module housing.
6. Close the controller module cover, as needed.

Step 6: Install the controller

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

> The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   > Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   > You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Interrupt the boot process **only** after determining the correct timing: You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. Only press `Ctrl-C` when you see the message `Press Ctrl-C for Boot Menu`. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. Select the option to boot to Maintenance mode from the displayed menu.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process. e. Interrupt the boot process **only** after determining the correct timing: You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. Only press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. f. From the boot menu, select the option for Maintenance mode.

Important: During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
  You can safely respond `y` to these prompts.

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the controller's HA state

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

- `ha`
- `mcc`
- `mcc-2n`
- `mccip`
- `non-ha`

b. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - FAS2700

To complete the replacement procedure and restore your system to full operation, you must recable the storage, confirm disk reassignment, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks.

You must use the correct procedure for your configuration:

Controller redundancy	Then use this procedure…
HA pair	Option 1: Verify the system ID change on an HA system
Stand-alone	Option 2: Manually reassign the system ID on a stand-alone system in ONTAP
Two-node MetroCluster configuration	Option 3: Manually reassign the system ID on systems in a two-node MetroCluster configuration

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *`local-node-name`* `partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

- Restore onboard key management encryption keys
- Restore external key management encryption keys

6. Give back the controller:

   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Option 2: Manually reassign the system ID on a stand-alone system in ONTAP

In a stand-alone system, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems that are in a stand-alone configuration.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by pressing Ctrl-C, and then select the option to boot to Maintenance mode from the displayed menu.
2. You must enter `Y` when prompted to override the system ID due to a system ID mismatch.
3. View the system IDs: `disk show -a`
4. You should make a note of the old system ID, which is displayed as part of the disk owner column.

   The following example shows the old system ID of 118073209:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                   POOL   SERIAL NUMBER  HOME
--------    -------------           -----  -------------  -------------
disk_name    system-1  (118073209)  Pool0  J8XJE9LC       system-1
(118073209)
disk_name    system-1  (118073209)  Pool0  J8Y478RC       system-1
(118073209)
.
.
.
```

5. Reassign disk ownership by using the system ID information obtained from the disk show command: `disk reassign -s old system ID disk reassign -s 118073209`
6. Verify that the disks were assigned correctly: `disk show -a`

   The disks belonging to the replacement node should show the new system ID. The following example now show the disks owned by system-1 the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                   POOL   SERIAL NUMBER  HOME
--------    -------------           -----  -------------  -------------
disk_name    system-1  (118065481)  Pool0  J8Y0TDZC       system-1
(118065481)
disk_name    system-1  (118065481)  Pool0  J8Y0TDZC       system-1
(118065481)
.
.
.
```

7. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

- Restore onboard key management encryption keys
- Restore external key management encryption keys

8. Boot the node: `boot_ontap`

Option 3: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `metrocluster node show -fields node-systemid,dr-partner-systemid`

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

   ```
   dr-group-id cluster         node                 node-systemid dr-
   partner-systemid
    ----------- --------------------- ------------------- -------------
   -------------------
    1           Cluster_A             Node_A_1             536872914
   118073209
    1           Cluster_B             Node_B_1             118073209
   536872914
    2 entries were displayed.
   ```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

   ```
   Local System ID: 118065481
       ...
       ...
   ```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                  POOL   SERIAL NUMBER  HOME
-------    -------------          -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Verify that the coredumps are saved: `system node run -node local-node-name partner savecore`

      If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node local-node-name partner savecore -s command`.</info>.

   c. Return to the admin privilege level: `set -privilege admin`

7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`

8. Boot the *replacement* node: `boot_ontap`

9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`

10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

       You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

    b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`
    c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - FAS2700

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.

4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS2700

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

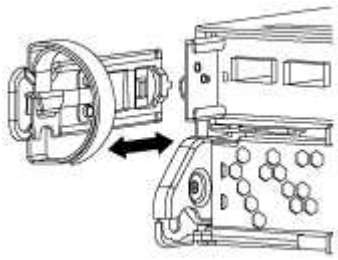

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

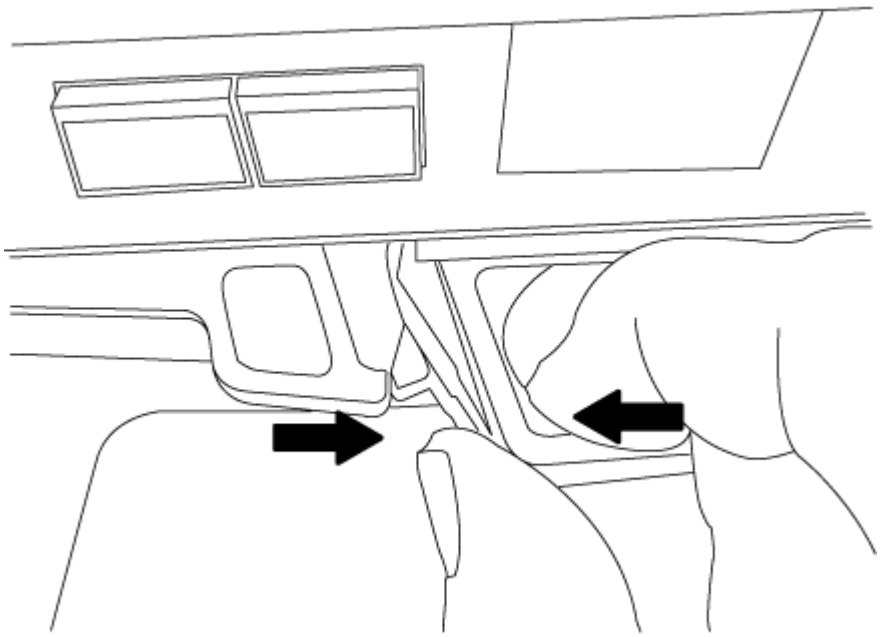

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

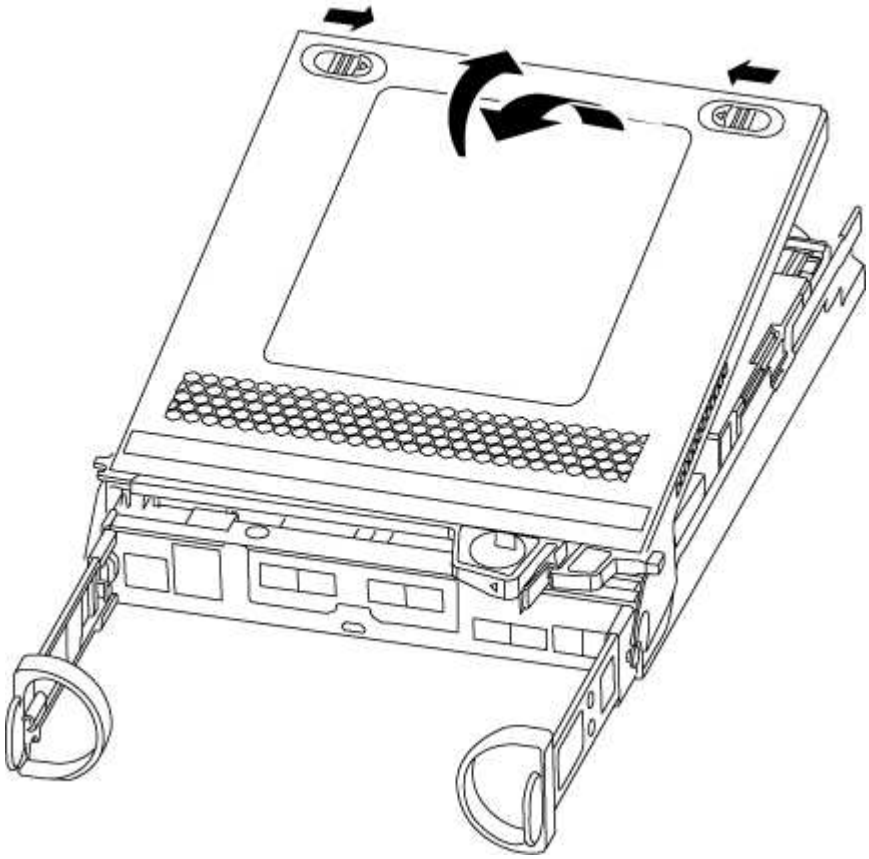

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

If you are replacing a DIMM, you need to remove it after you have unplugged the NVMEM battery from the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED on the back of controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

   

3. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
4. If the NVMEM LED is flashing, there is data in the NVMEM and you must disconnect the battery to clear the memory:
   a. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

   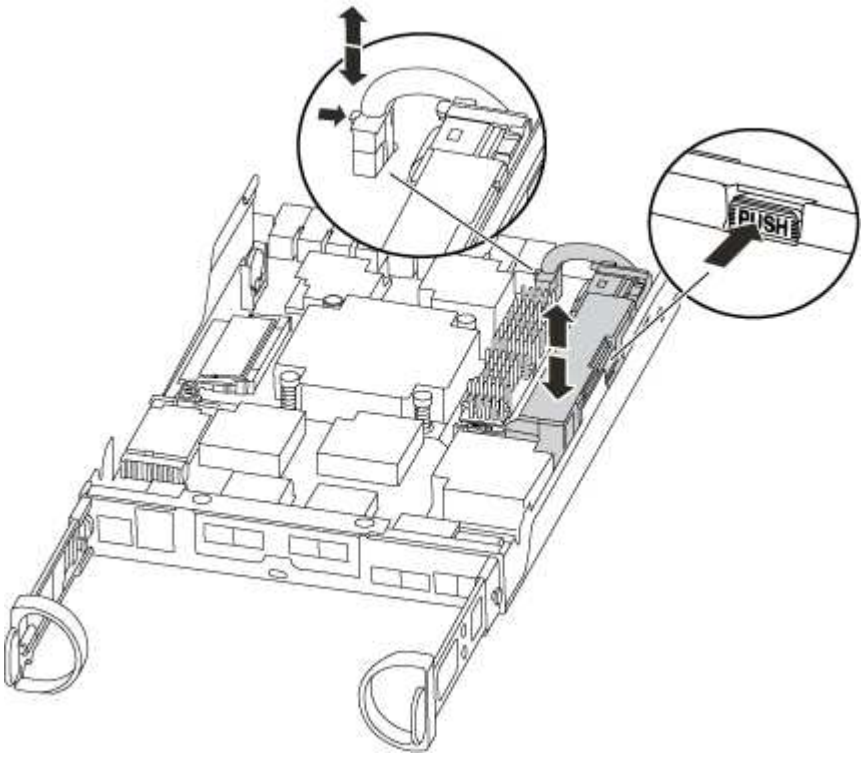

   b. Confirm that the NVMEM LED is no longer lit.
   c. Reconnect the battery connector.

5. Return to Step 3: Replace the DIMMs in this procedure to recheck the NVMEM LED.
6. Locate the DIMMs on your controller module.
7. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
8. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   The number and placement of system DIMMs depends on the model of your system.

   The following illustration shows the location of system DIMMs:

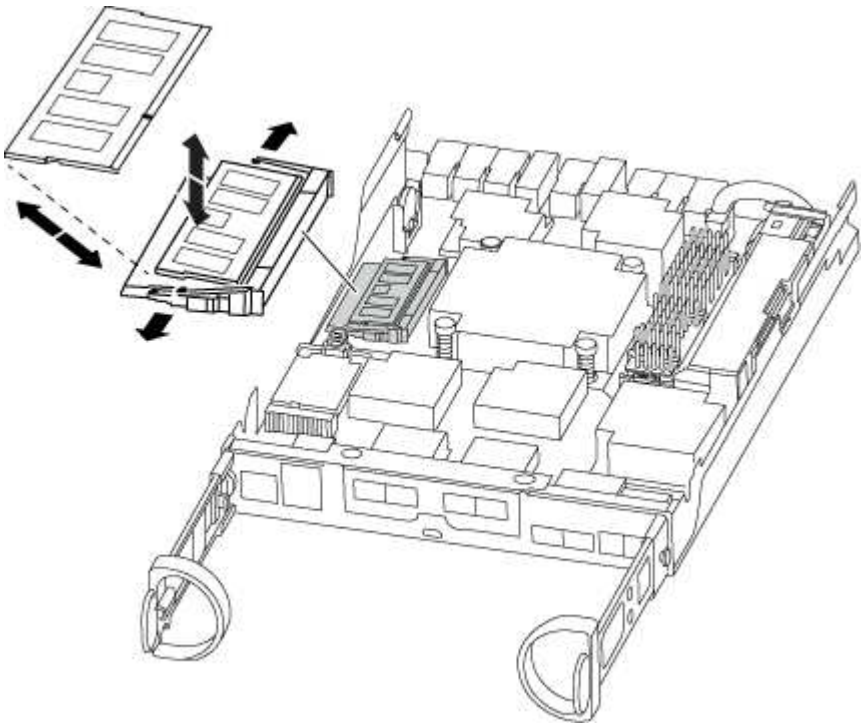

9. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

10. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

    The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

    Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

11. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
12. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

Make sure that the plug locks down onto the controller module.

13. Close the controller module cover.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - AFF A220 and FAS2700

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace the NVMEM battery - FAS2700

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

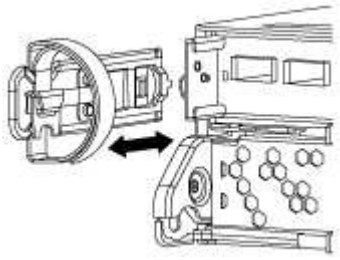

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

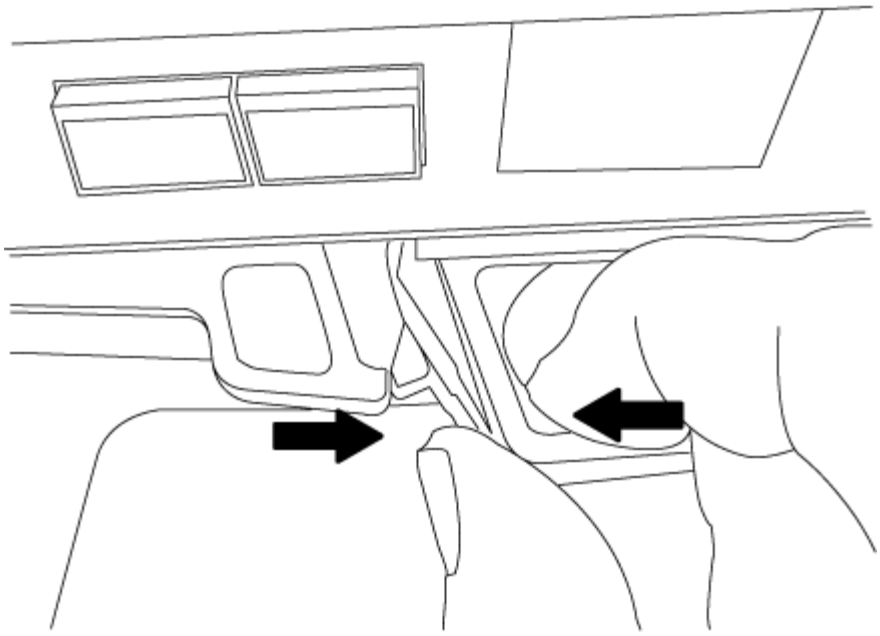

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery in your system, you must remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

3. Locate the NVMEM battery in the controller module.

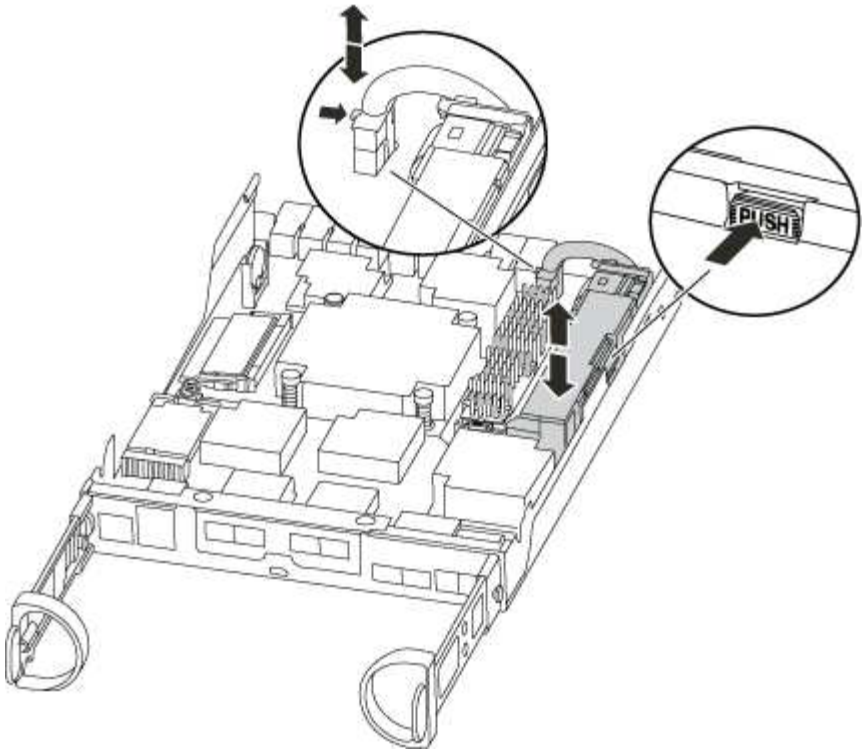

4. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
5. Remove the battery from the controller module and set it aside.
6. Remove the replacement battery from its package.
7. Loop the battery cable around the cable channel on the side of the battery holder.
8. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
9. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
10. Plug the battery plug back into the controller module.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber

optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - FAS2700

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> Cooling is integrated with the power supply, so you must replace the power supply within two minutes of removal to prevent overheating due to reduced airflow. Because the chassis provides a shared cooling configuration for the two HA nodes, a delay longer than two minutes will shut down all controller modules in the chassis. If both controller modules do shut down, make sure that both power supplies are inserted, turn both off for 30 seconds, and then turn both on.

- Power supplies are auto-ranging.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.

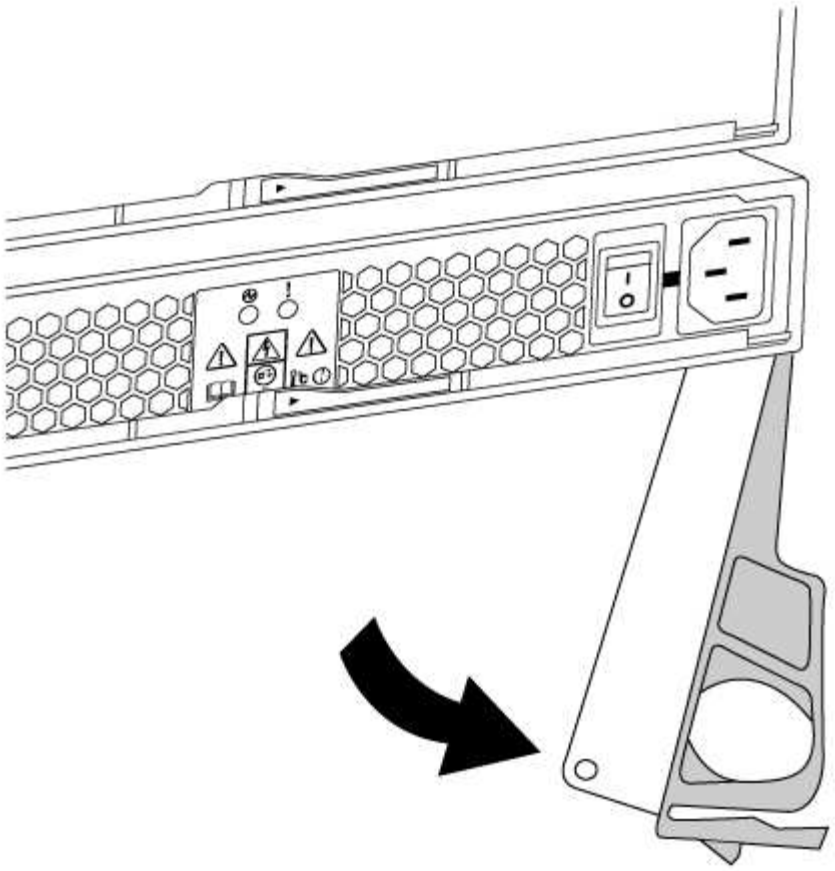

5. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.
7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
9. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.

10. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

    The power supply LEDs are lit when the power supply comes online.

11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - FAS2700

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

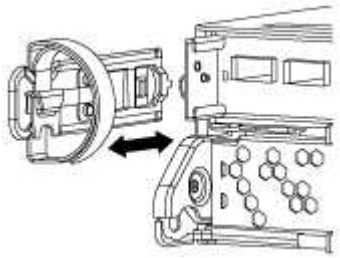

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

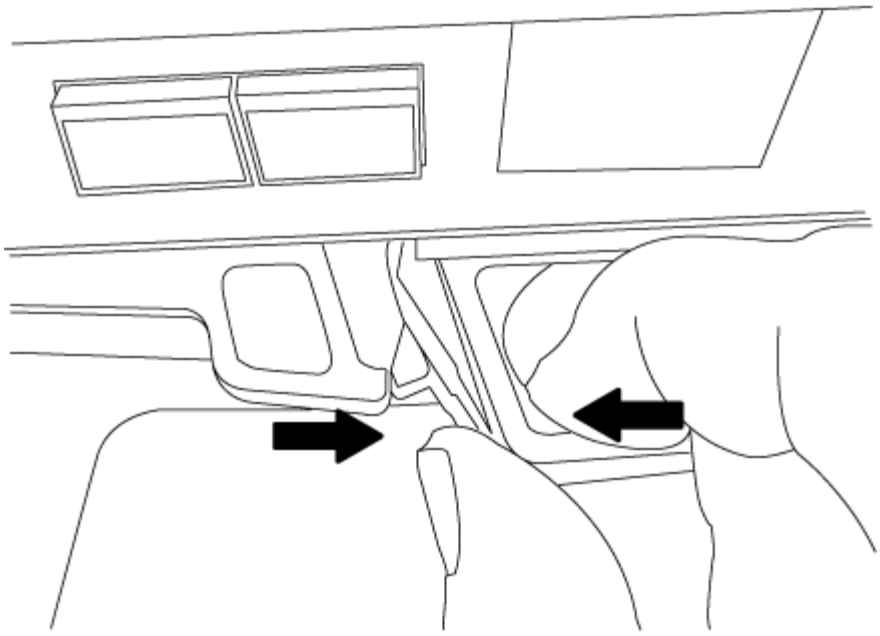

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

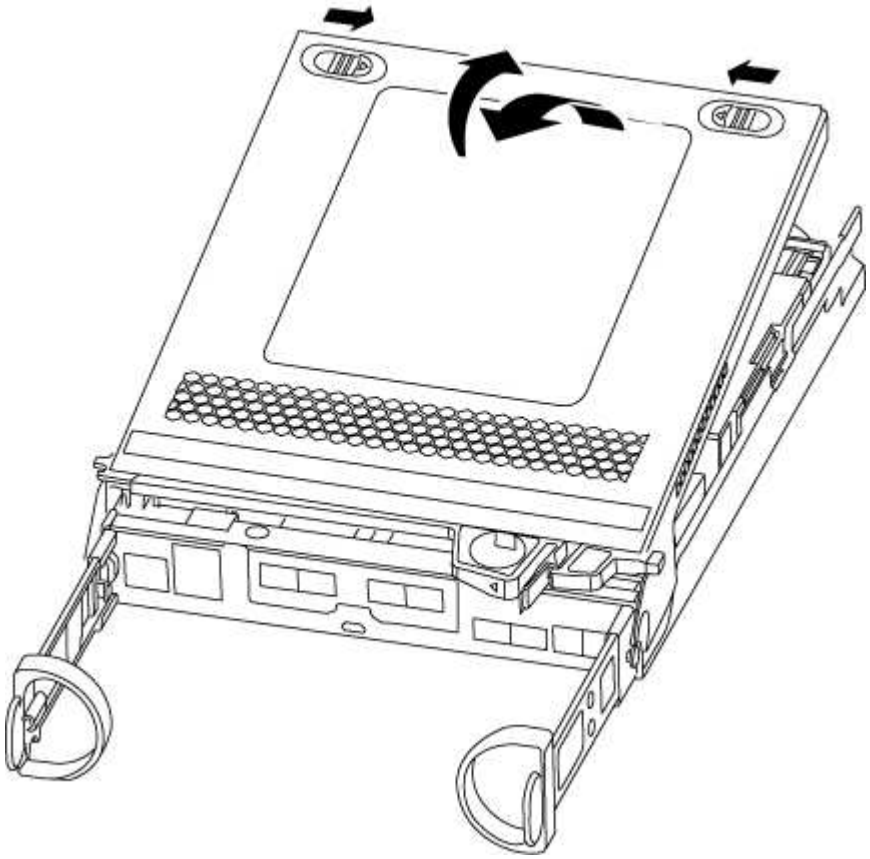

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

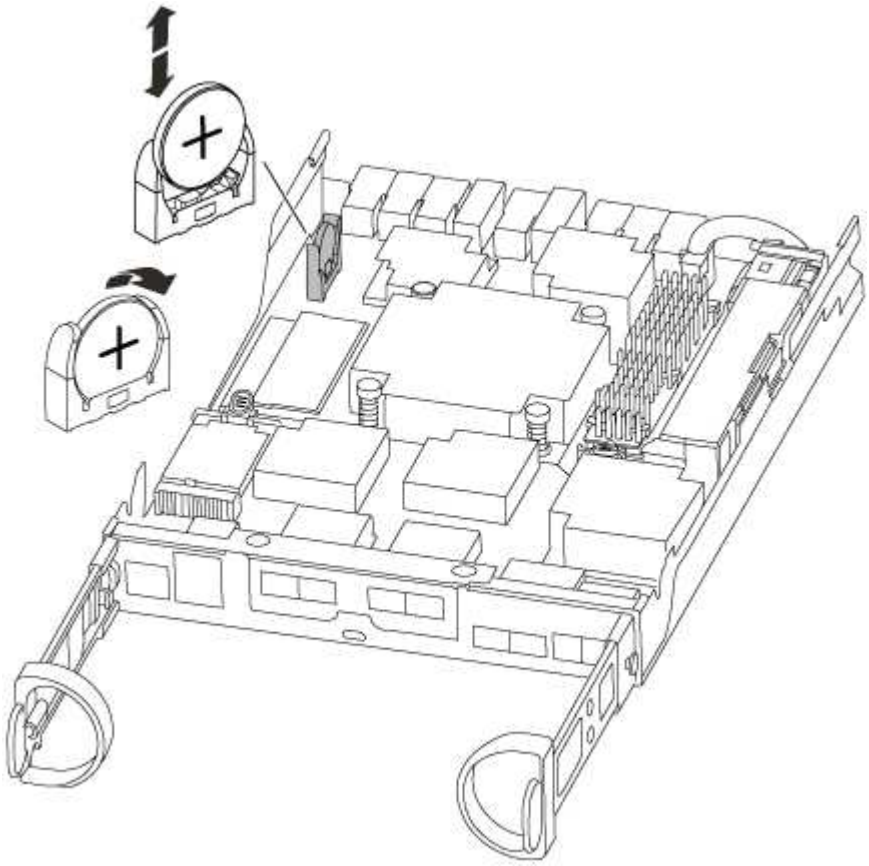

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and set time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

FAS2820 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick guide - FAS2820

This guide gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Access the *Installation and Setup Instructions* PDF poster:

FAS2820 Systems Installation and Setup Instructions

Video steps - FAS2820

The following video shows how to install and cable your new system.

Animation - FAS2820 Installation and setup instructions

Detailed steps - FAS2820

This procedure gives detailed step-by-step instructions for installing a typical NetApp storage system. Use this procedure if you want more detailed installation instructions.

Step 1: Prepare for installation

Before you begin

You need to provide the following at your site:

- Rack space for the storage system in either a telco rack or system cabinet.
  - 2U for the storage system
  - 2U or 4U for each drive shelf in your system
- Phillips #2 screwdriver
- Additional networking cables to connect your storage system to your network switch and laptop or console with a Web browser
- A laptop or console with an RJ-45 connection and access to a Web browser
  - Access to the NetApp Hardware Universe for information about site requirements as well as additional information on your configured storage system.
  - You might also want to have access to the Release Notes for your version of ONTAP 9 for your version of ONTAP for more information about this storage system.

Steps

1. Unpack all boxes and inventory the contents.

   Customers with specific power requirements must check NetApp Hardware Universe for their configuration options.

2. Access the Configure ONTAP on a new cluster with System Manager
   a. Review the requirements amd procedure steps.
   b. Gather information about your storage system by completing the setup worksheet^ (need the URL to the worksheet).
   c. Record the storage system serial number from the controllers.

The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
10 GbE, SFP28 cable (order dependent)	X6566B-05-R6, .5, X6566B-2-R6, 2m		Network cable
25Gb Ethernet, SFP28	X66240A-05, .5m X66240-2, 2m X66240A-5, 5m		Network cable

Type of cable…	Part number and length	Connector type	For…
32Gb Fiber Channel, SFP+ (target/initiator)	X66250-2, 2m X66250-5, 5m X66250-15, 15m		FC network
Cat 6, RJ-45 (order dependent)	X6561-R6 X6562-R6		Management network and Ethernet data
Storage	X66030A, 0.5m X66031A, 1m X66032A, 2m		Storage
USB-C console cable	No part number label		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	No part number label		Powering up the storage system
Optional FC cable	Optional FC cable		Additional FC network cable

Step 2: Install the hardware

Install your storage system in a telco rack or NetApp storage system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your storage system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the storage system.

   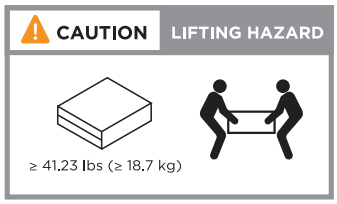

3. Attach cable management devices (as shown).

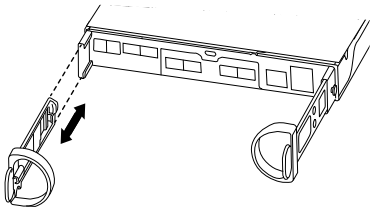

4. Place the bezel on the front of the storage system.

Step 3: Cable controllers to your network

Cable the controllers to your network as either a two-node switchless cluster or a switched cluster.

The following table identifies the cable type with the call out number and cable color in the illustrations for both two-node switchless cluster and switched cluster.

Cabling	Connection type
1	Cluster interconnect
2	Management network switch
3	Host network switches

Before you begin

- Contact your network administrator for information about connecting the storage system to the switches.
- Check the illustration arrow for the proper cable connector pull-tab orientation.
  - As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn the cable head over and try again.
  - If connecting to an optical switch, insert the SFP into the controller port before cabling to the port.

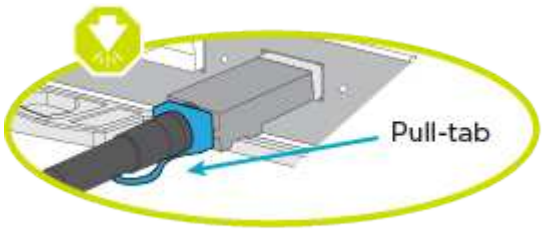

You can use either the applicable animation or detailed steps in the table to cable your controllers to your network.

Animation - Cabling a two-node switchless cluster cabling

Animation - Switched cluster cabling

Option 1: Cable a two-node switchless cluster

Cable your network connections and your cluster interconnect ports for a two-node switchless cluster.

Steps

1. Cable the cluster interconnect ports e0a to e0a and e0b to e0b with the cluster interconnect cable:

Cluster interconnect cables

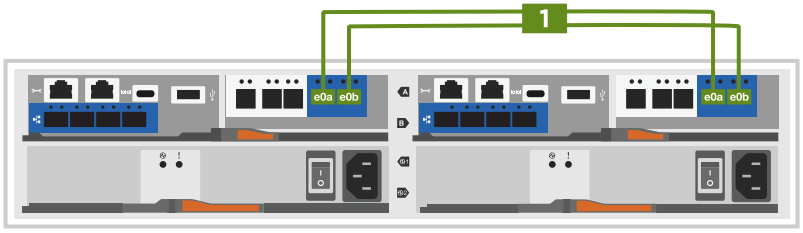

2. Cable the e0M ports to the management network switches with the RJ45 cables:

RJ45 cables

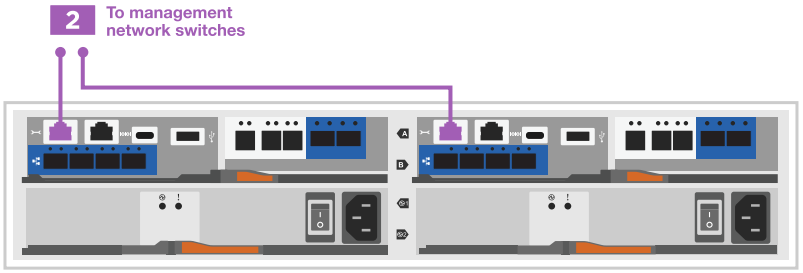

3. Cable the mezzanine card ports to your host network.

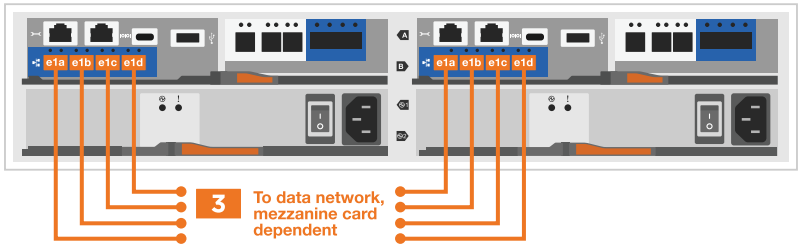

   a. If you have a 4-port Ethernet data network, cable ports e1a through e1d to your Ethernet data network.

      - 4-ports, 10/25Gb Ethernet, SFP28

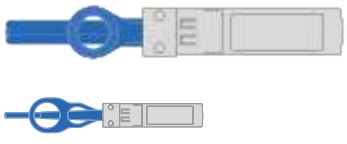

- 4-ports, 10GBASE-T, RJ45

b. If you have a 4-port Fiber Channel data network, cable ports 1a through 1d for your FC network.

- 4-ports, 32Gb Fiber Channel, SFP+ (target only)

- 4-ports, 32Gb Fiber Channel, SFP+ (initiator/target)

c. If you have a 2+2 card (2 ports with Ethernet connections and 2 ports with Fiber Channel connections), cable ports e1a and e1b to your FC data network and ports e1c and e1d to your Ethernet data network.

- 2-ports, 10/25Gb Ethernet (SFP28) + 2-ports 32Gb FC (SFP+)

DO NOT plug in the power cords.

Option 2: Cable a switched cluster

Cable your network connections and your cluster interconnect ports for a switched cluster.

Steps

1. Cable the cluster interconnect ports e0a to e0a and e0b to e0b with the cluster interconnect cable:

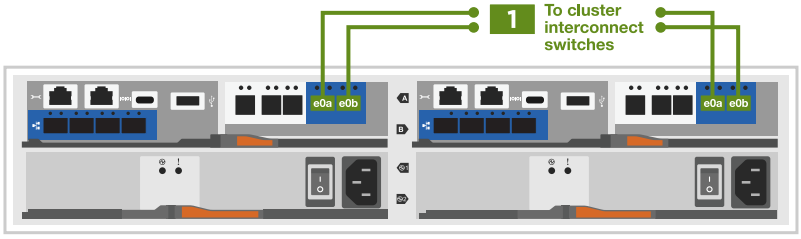

2. Cable the e0M ports to the management network switches with the RJ45 cables:

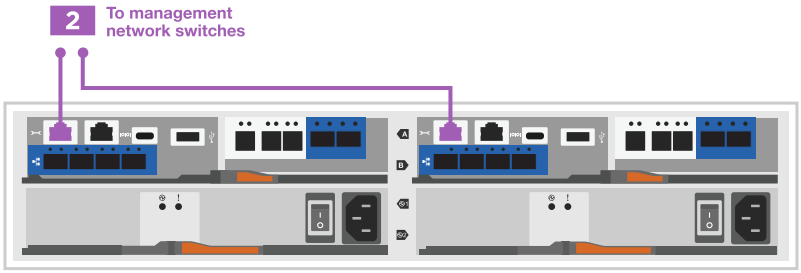

3. Cable the mezzanine card ports to your host network.

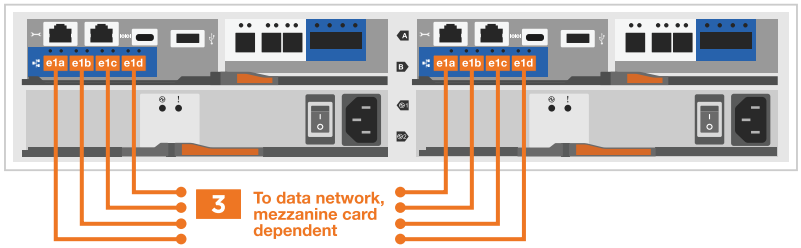

a. If you have a 4-port Ethernet data network, cable ports e1a through e1d to your Ethernet data network.

   - 4-ports, 10/25Gb Ethernet, SFP28

   - 4-ports, 10GBASE-T, RJ45

b. If you have a 4-port Fiber Channel data network, cable ports 1a through 1d for your FC network.

   - 4-ports, 32Gb Fiber Channel, SFP+ (target only)

   - 4-ports, 32Gb Fiber Channel, SFP+ (initiator/target)

c. If you have a 2+2 card (2 ports with Ethernet connections and 2 ports with Fiber Channel connections), cable ports e1a and e1b to your FC data network and ports e1c and e1d to your Ethernet data network.

   - 2-ports, 10/25Gb Ethernet (SFP28) + 2-ports 32Gb FC (SFP+)

DO NOT plug in the power cords.

Step 4: Cable controllers to drive shelves

Cable your controllers to external storage.

The following table identifies the cable type with the call out number and cable color in the illustrations for cabling your drive shelves to your storage system.

The example uses DS224C. Cabling is similar with other supported drive shelves. See Install and cable shelves for a new system installation - shelves with IOM12/IOM12B modules for more information.

Cabling	Connection type
1	Shelf-to-shelf cabling
2	Controller A to the drive shelves
3	Controller B to the drive shelves

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

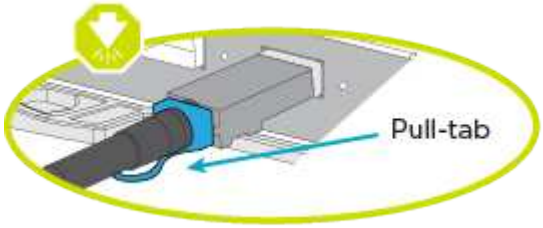

About this task

Use the animation or the step-by step instructions to complete the cabling between the controllers and to the drive shelves.

Do not use port 0b2 on a FAS2820. This SAS port is not used by ONTAP and is always disabled. See Install a shelf in a new storage system for more information.

Animation - Drive shelf cabling

Steps

1. Cable the shelf-to-shelf ports.
   a. Port 1 on IOM A to port 3 on the IOM A on the shelf directly below.
   b. Port 1 on IOM B to port 3 on the IOM B on the shelf directly below.

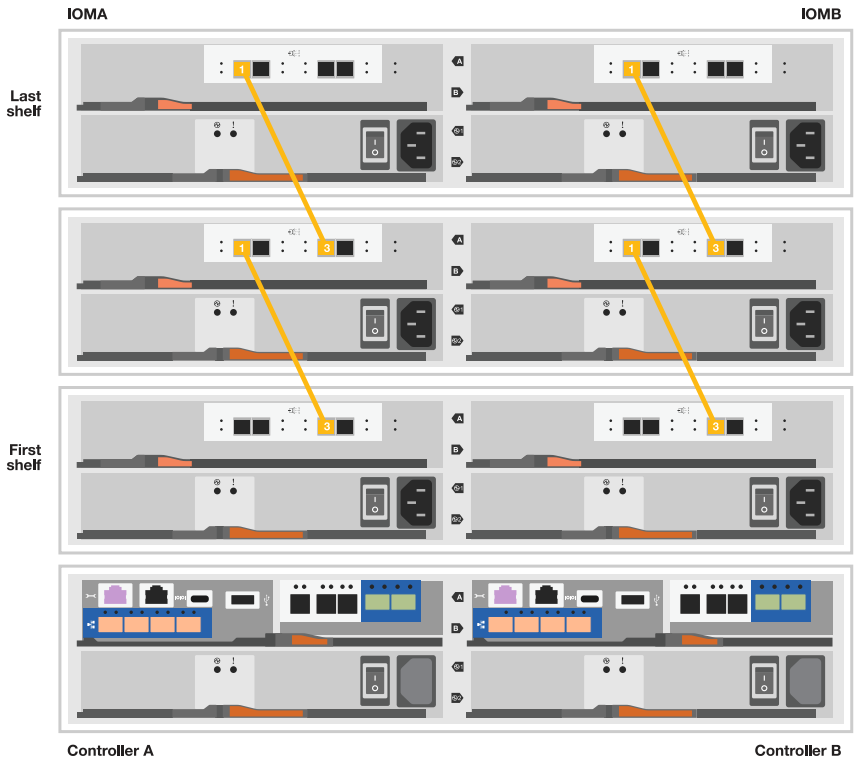

2. Cable controller A to the drive shelves.
   a. Controller A port 0a to IOM B port 1 on first drive shelf in the stack.
   b. Controller A port 0b1 to IOM A port 3 on the last drive shelf in the stack.

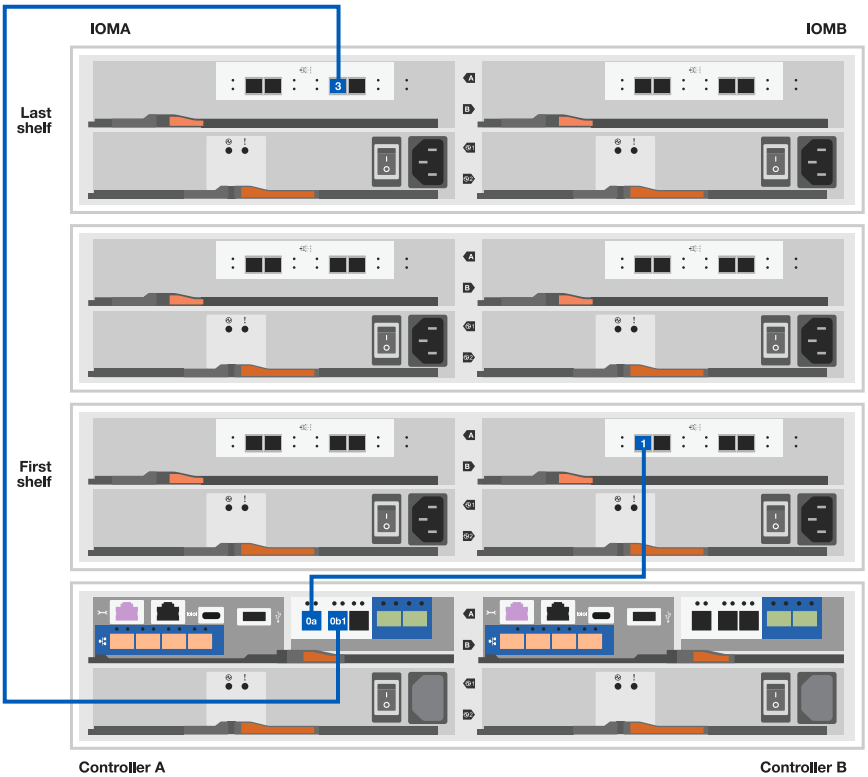

3. Connect controller B to the drive shelves.

a. Controller B port 0a to IOM A port 1 on first drive shelf in the stack.

b. Controller B port 0b1 to IOM B port 3 on the last drive shelf in the stack.

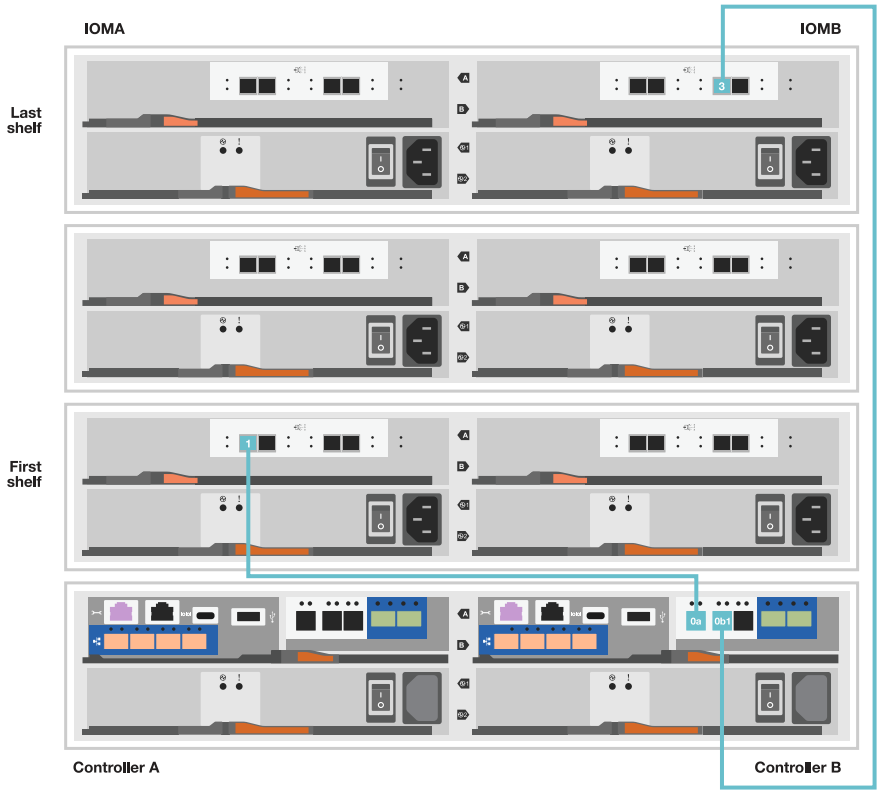

Step 5: Complete storage system setup and configuration

Complete your storage system setup and configuration using either Option 1: if network discovery enabled or Option 2: if network discovery is not enabled.

Use the following animation in either option where setting shelf ID is required:

Animation - Set drive shelf IDs

Option 1: If network discovery is enabled

If network discovery enabled on your laptop, complete storage system setup and configuration using automatic cluster discovery.

Steps

1. Turn on shelf power and set shelf IDsusing the animation at the beginning of this Step.
2. Power on the controllers
   a. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
   b. Turn on the power switches to both nodes.

      Initial booting may take up to eight minutes.

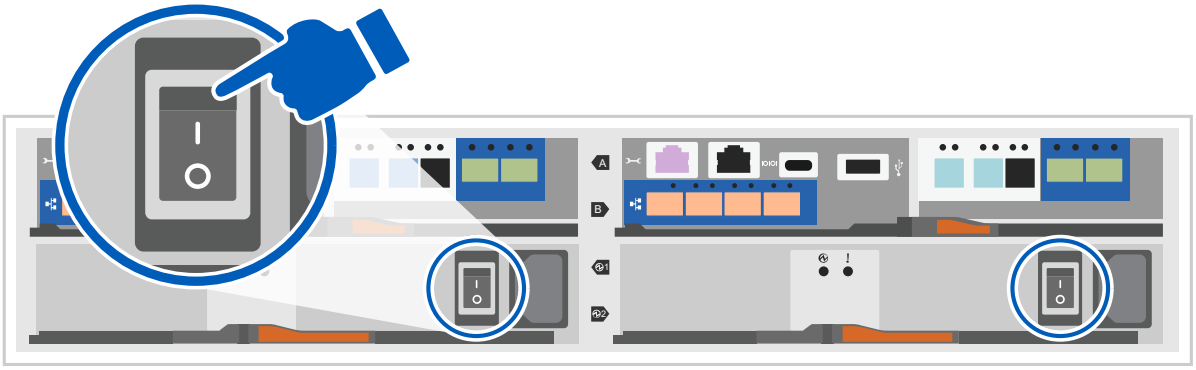

3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Connect your laptop to the Management switch.
5. Use the graphic or steps to discover the storage system node to configure::

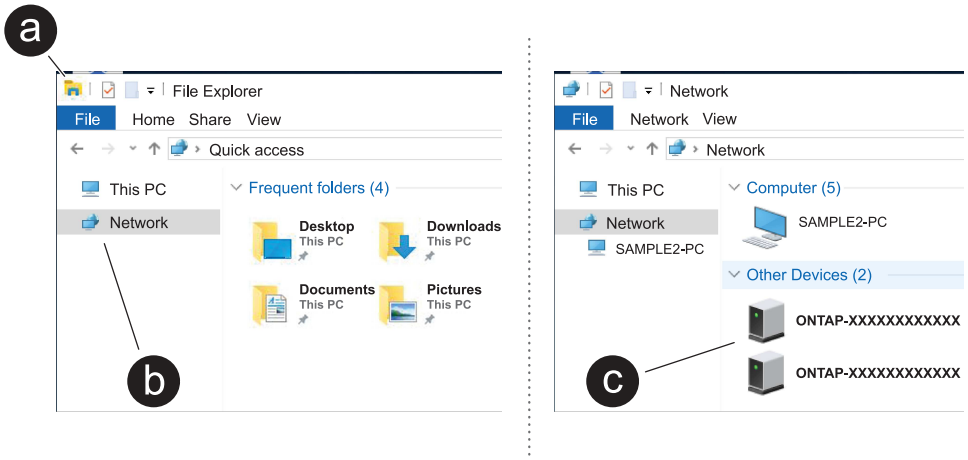

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

XXXXX is the storage system serial number for the target node.

System Manager opens.

6. Use System Manager guided setup to configure your storage system using the data you collected in Step 1: Prepare for installation.
7. Create an account or log into your account.
   a. Click mysupport.netapp.com
   b. Click *Create Account* if you need to create an account or log into your account.
8. Download and install Active IQ Config advisor
   a. Verify the health of your storage system by running Active IQ Config Advisor.
9. Register your system at https://mysupport.netapp.com/site/systems/register.
10. After you have completed the initial configuration, go to the NetApp ONTAP Resources page for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

If network discovery is not enabled on your laptop, manually complete the configuration and setup.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your storage system, and then connect the laptop or console to the switch on the management subnet.

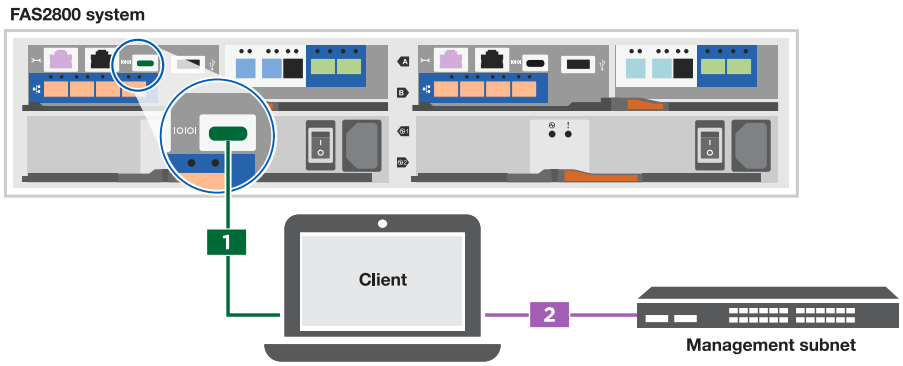

   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Turn on shelf power and set shelf IDs using the animation at the beginning of this Step.
3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
4. Turn on the power switches to both nodes.

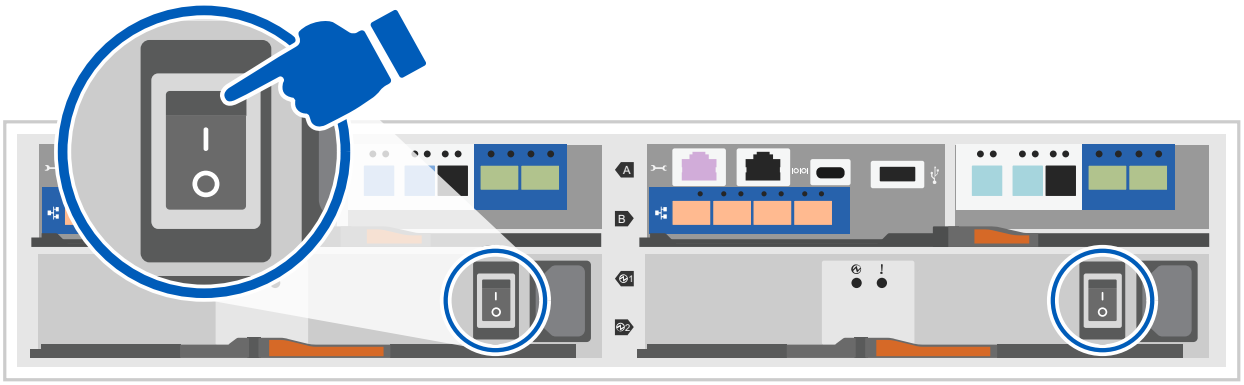

Initial booting may take up to eight minutes.

5. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

6. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.
   The format for the address is https://x.x.x.x.
   b. Configure the storage system using the data you collected in Step 1: Prepare for installation..
7. Create an account or log into your account.
   a. Click mysupport.netapp.com
   b. Click *Create Account* if you need to create an account or log into your account.
8. Download and install Active IQ Config advisor
   a. Verify the health of your storage system by running Active IQ Config Advisor.
9. Register your system at https://mysupport.netapp.com/site/systems/register.
10. After you have completed the initial configuration, go to the NetApp ONTAP Resources page for information about configuring additional features in ONTAP.

Maintain

Maintain FAS2820 hardware

Maintain the hardware of your FAS2820 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS2820 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS2820 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Caching module	You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Drive	A drive is a device that provides the physical storage media for data.
NVMEM battery	A battery is included with the controller and preserves cached data if the AC power fails.
Mezzanine card	A Mezzanine card is an expansion card that is designed to be inserted into a specialized slot on the motherboard and holds the card I/O cards.

Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - FAS2800

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your FAS2800 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the System Management module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - FAS2800

Before replacing the boot media in your FAS2800 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - FAS2800

Shut down the impaired controller in your FAS2800 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - FAS2800

The boot media in your FAS2800 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module and is accessed by removing the controller module from the chassis and removing the controller module cover.module from the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the

system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

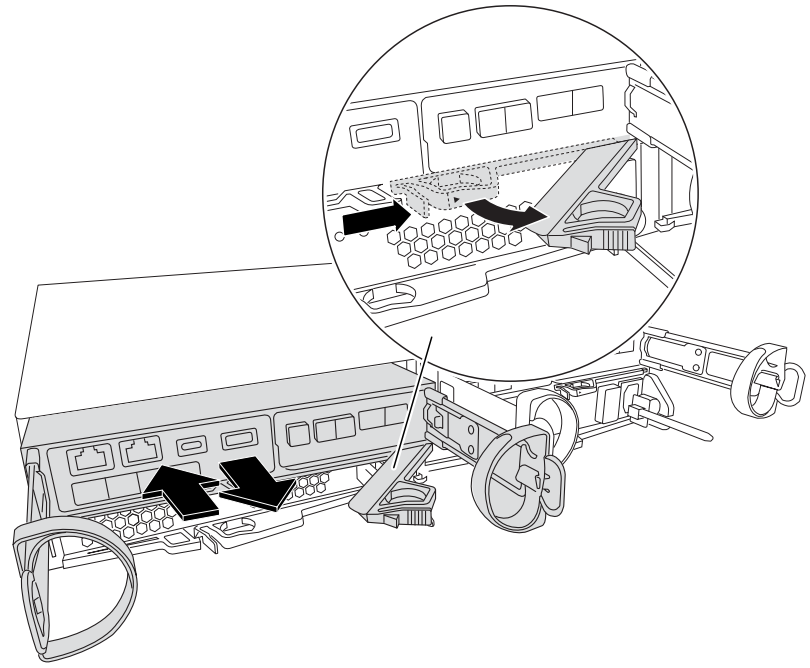

4. Turn the controller module over and place it on a flat, stable surface.
5. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.

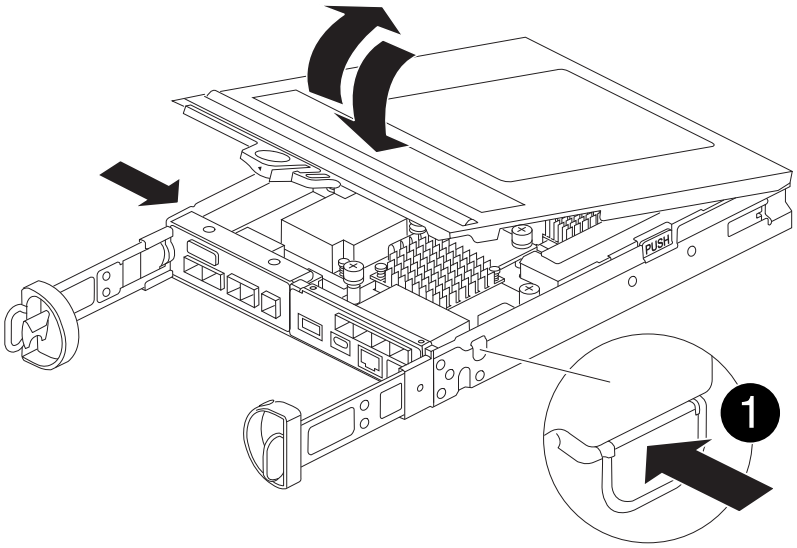

1	Controller module cover release button

6. Locate the boot media in the controller module, located under the mezzanine card and follow the directions to replace it.

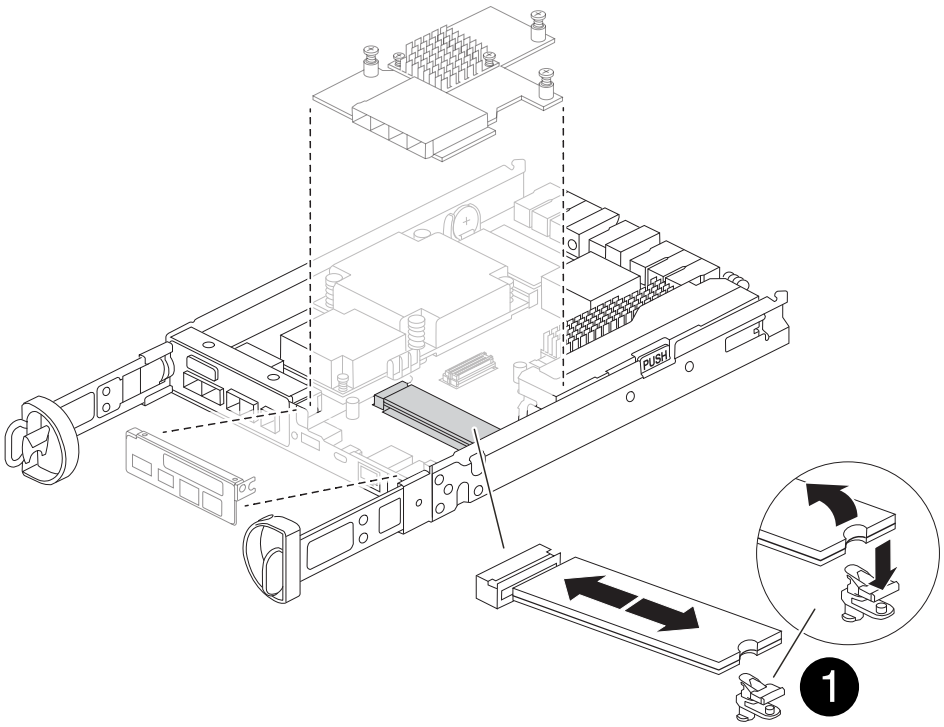

1	Boot media locking tab

7. Remove the mezzanine card using the following illustration or the FRU map on the controller module:
   a. Remove the IO Plate by sliding it straight out from thr controller module.
   b. Loosen the thumbscrews on the mezzanine card.

      You can loosen the thumbscrews with your fingers or a screwdriver. If you use your fingers, you might need to rotate the NV battery up for better finger purchase on the thumbscrew next to it.

   c. Lift the mezzanine card straight up.
8. Replace the boot media:
   a. Press the blue button on the boot media housing to release the boot media from its housing, rotate the boot media up, and then gently pull it straight out of the boot media socket.

      Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

   b. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
      Check the boot media to make sure that it is seated squarely and completely in the socket, and if necessary, remove the boot media and reseat it into the socket.
   c. Push the blue locking button, rotate the boot media all the way down, and then release the locking button to lock the boot media in place.
9. Reinstall the mezzanine card:
   a. Align the socket on the motherboard with the socket on the mezzanine card, and then gently seat the card in the socket.

b. Tighten the three thumbscrews on the mezzanine card.

c. Reinstall the IO Plate.

10. Reinstall the controller module cover and lock it into place.

11. Install the controller module:

    a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module half-way into the way into the system.

    b. Recable the controller, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

       The controller begins to boot and stops at the LOADER prompt.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - FAS2800

After installing the new boot media device in your FAS2800 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - FAS2800

If a component in your FAS2800 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot Media - manual recovery

Boot media manual recovery workflow - FAS2800

Get started with replacing the boot media in your FAS2800 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - AFF A800

Before replacing the boot media in your AFF A800 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - FAS2820

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

What's next?

After checking the encryption key support and status on the boot media, you need to shut down the controller.

Shut down the controller for manual boot media recovery - FAS2820

Shut down or take over the impaired controller.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

What's next?

After shutting down the controller, you need to replace the boot media.

Replace the boot media and prepare for manual boot recovery - FAS2820

You must remove and open the impaired controller module, locate and replace the boot media in the controller, transfer the boot image to a USB drive, insert the USB drive in the controller, and then boot the controller.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

Steps

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.
3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

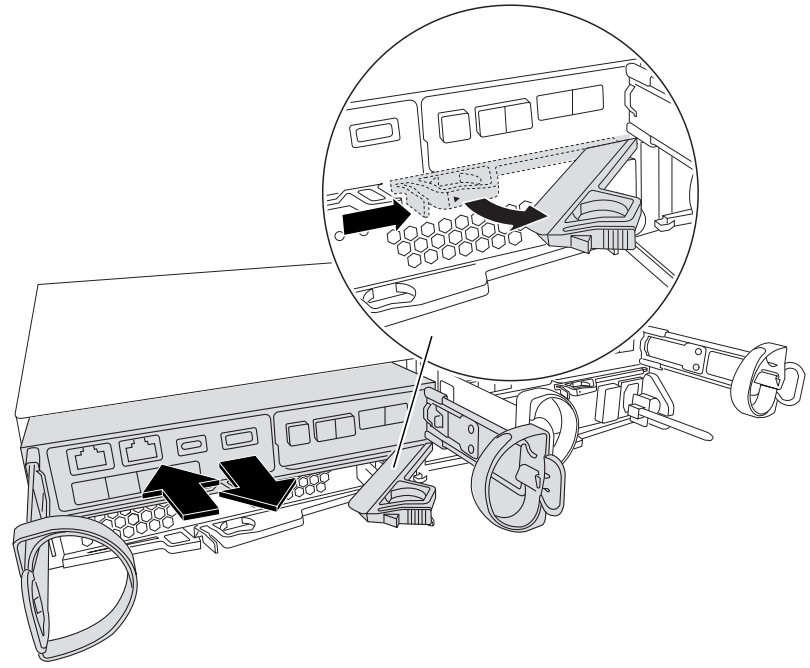

4. Turn the controller module over and place it on a flat, stable surface.
5. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.

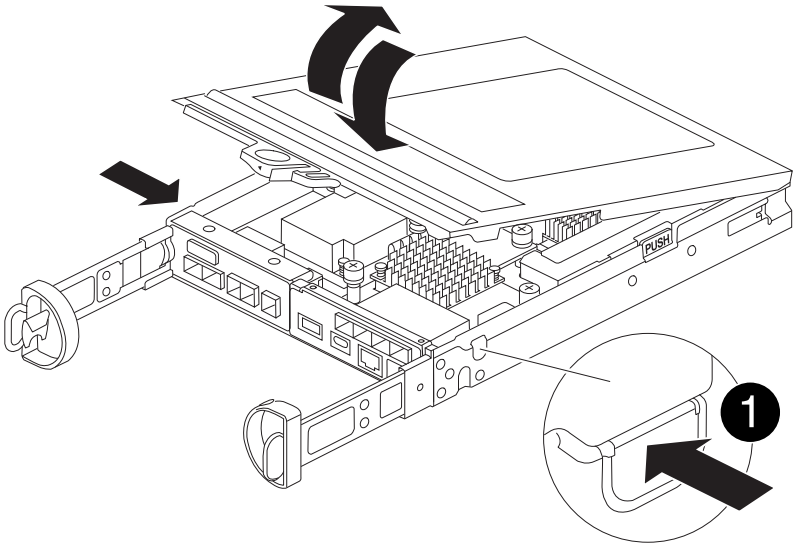

1	Controller module cover release button

Step 2: Replace the boot media

Locate the boot media in the controller module, located under the mezzanine card and follow the directions to replace it.

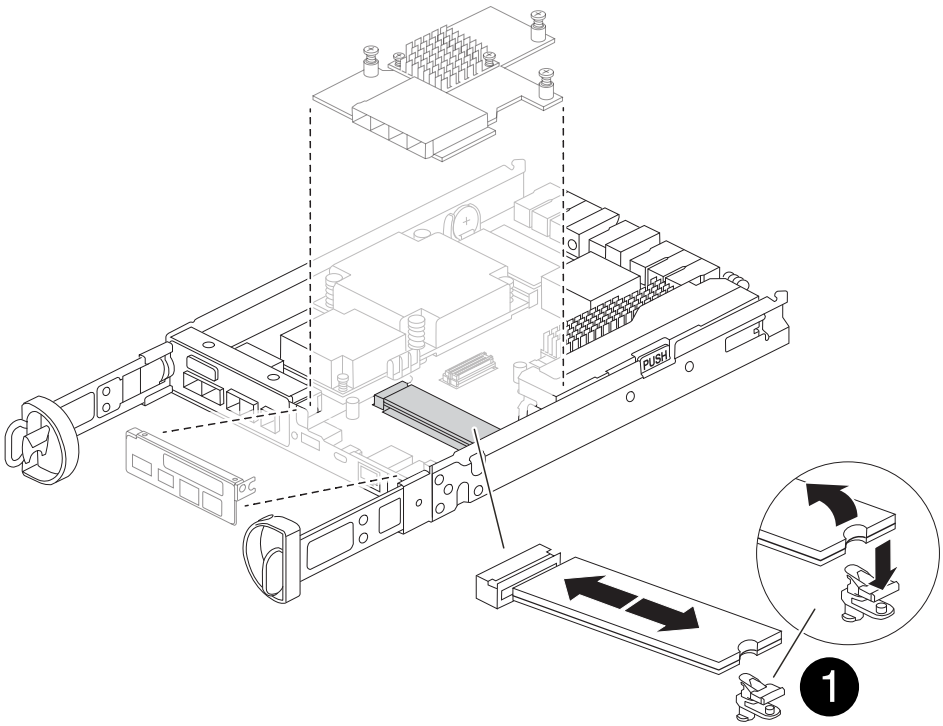

1	Boot media locking tab

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the mezzanine card using the following illustration or the FRU map on the controller module:
   a. Remove the IO Plate by sliding it straight out from thr controller module.
   b. Loosen the thumbscrews on the mezzanine card.

      You can loosen the thumbscrews with your fingers or a screwdriver. If you use your fingers, you might need to rotate the NV battery up for better finger purchase on the thumbscrew next to it.

   c. Lift the mezzanine card straight up.
3. Replace the boot media:
   a. Press the blue button on the boot media housing to release the boot media from its housing, rotate the boot media up, and then gently pull it straight out of the boot media socket.

      Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

   b. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
      Check the boot media to make sure that it is seated squarely and completely in the socket, and if necessary, remove the boot media and reseat it into the socket.
   c. Push the blue locking button, rotate the boot media all the way down, and then release the locking button to lock the boot media in place.

4. Reinstall the mezzanine card:
    a. Align the socket on the motherboard with the socket on the mezzanine card, and then gently seat the card in the socket.
    b. Tighten the three thumbscrews on the mezzanine card.
    c. Reinstall the IO Plate.
5. Reinstall the controller module cover and lock it into place.

Step 3: Transfer the boot image to the boot media

Install the system image on the replacement boot media using a USB flash drive with the image installed on it. You must restore the var file system during this procedure.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity.
- You must have a network connection.

Steps

1. Download the appropriate image version of ONTAP to the formatted USB flash drive:
    a. Use How to determine if the running ONTAP version supports NetApp Volume Encryption (NVE) to determine if volume encryption is currently supported.
        - If NVE is supported on the cluster, download the image with NetApp Volume Encryption.
        - If NVE is not supported on the cluster, download the image without NetApp Volume Encryption. See Which ONTAP image should I download? With or without Volume Encryption? for more details.
2. Remove the USB flash drive from your laptop.
3. Install the controller module:
    a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
    b. Recable the controller module.

    When recabling, remember to reinstall the media converters (SFPs) if they were removed.
4. Insert the USB flash drive into the USB slot on the controller module.

    Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.
5. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

    The controller begins to boot as soon as it is completely installed into the chassis and stops at the LOADER prompt.

What's next?

After replacing the boot media, you need to boot the recovery image.

Manual boot media recovery from a USB drive - FAS2820

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `...would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

What's next?

After booting the recovery image, you need to restore encryption on the boot media.

Restore encryption - FAS2820

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    `security key-manager key query -restored false.`

    ```
    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    `storage failover giveback -fromnode local`

13. Restore automatic giveback, if you disabled it, by entering the following command.

    `storage failover modify -node local -auto-giveback true`

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    `system node autosupport invoke -node * -type all -message MAINT=END`

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
**********************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
**********************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

What's next?

After restoring encryption on the boot media, you need to return the failed part to NetApp.

Return the failed boot media to NetApp - FAS2820

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the caching module - FAS2820

You must replace the caching module in the controller module when your system registers a single AutoSupport (ASUP) message that the module has gone offline; failure to do so results in performance degradation.

- You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller.

ONTAP 9 System Administration Reference

You might want to erase the contents of your caching module before replacing it.

Steps

1. Although data on the caching module is encrypted, you might want to erase any data from the impaired caching module and verify that the caching module has no data:

   a. Erase the data on the caching module: `system controller flash-cache secure-erase run -node node name localhost -device-id device_number`

      Run the `system controller flash-cache show` command if you don't know the Flash Cache device ID.

   b. Verify that the data has been erased from the caching module: `system controller flash-cache secure-erase show`

      The output should display the caching module status as erased.

2. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

3. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

4. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller: • For an HA pair, take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`. • For a stand-alone system: `system node halt` *`impaired_node_name`*

Step 2: Remove controller module

Remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
3. Remove and set aside the cable management devices from the left and right sides of the controller module.
4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

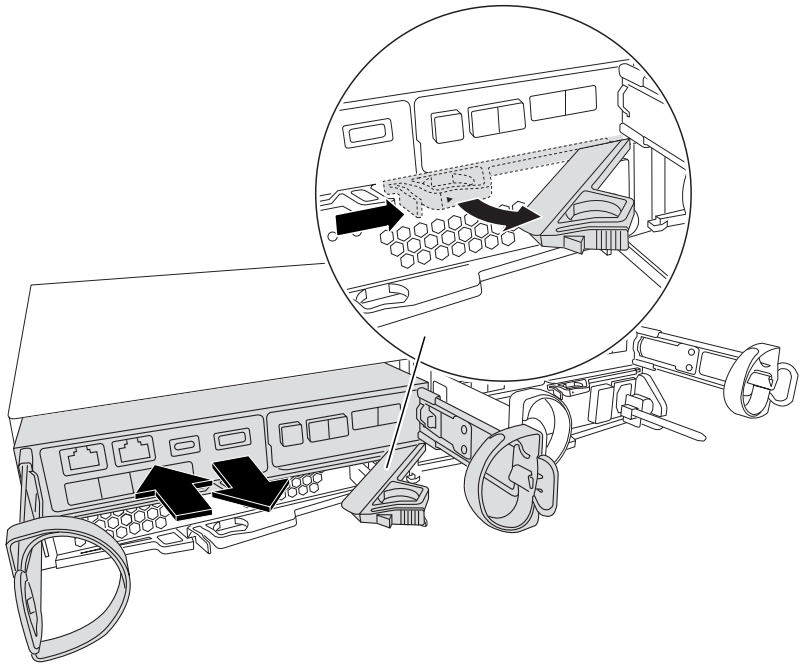

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.

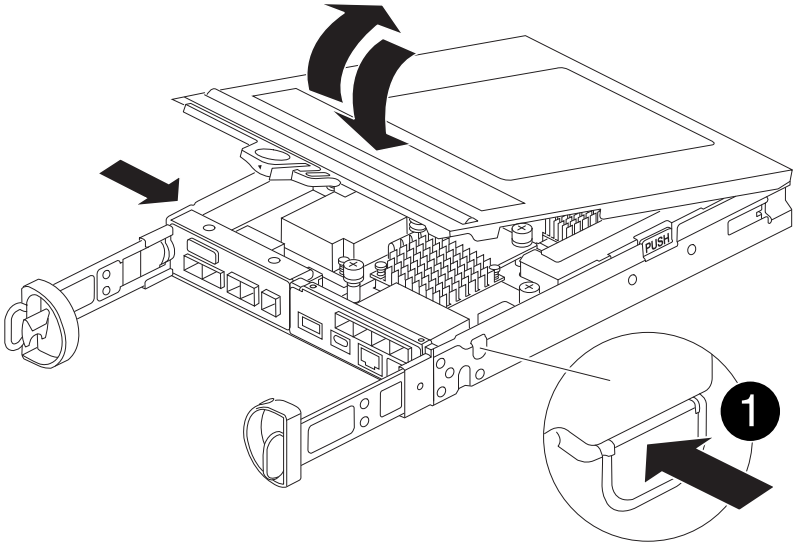

1	Controller module cover release button

Step 3: Replace a caching module

Locate the caching module inside the controller, remove the failed caching module and replace it.

Animation - Replace the caching module

Your storage system must meet certain criteria depending on your situation:

- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- All other components in the storage system must be functioning properly; if not, you must contact technical support.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed caching module near the rear of the controller module and remove it.
   a. Press the blue release button and rotate the caching module upward.
   b. Gently pull the caching module straight out of the housing.

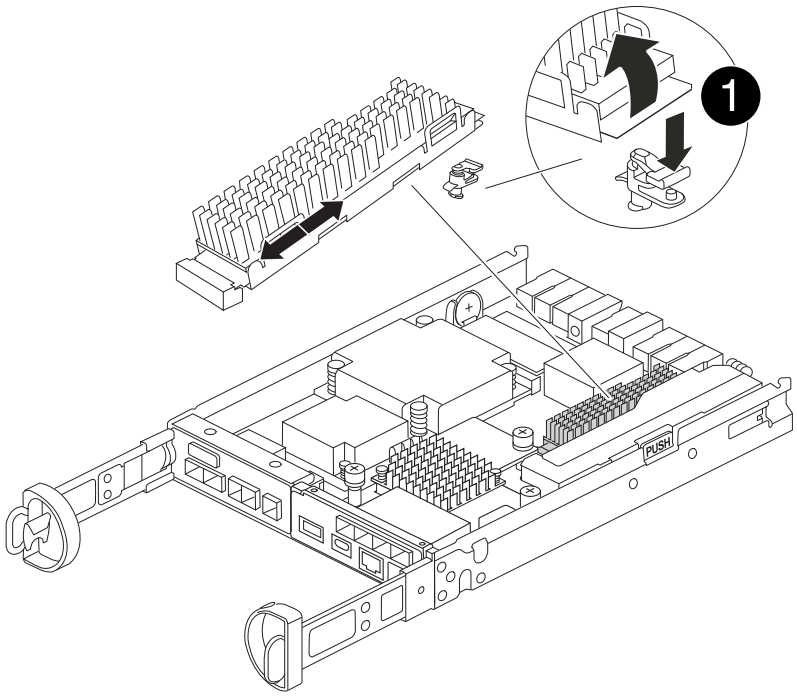

1	Caching module release button

3. Align the edges of the replacement caching module with the socket in the housing, and then gently push it into the socket.
4. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

5. Push the blue locking button, rotate the caching module all the way down, and then release the locking button to lock the caching module in place.
6. Reinstall the controller module cover and lock it into place.

Step 4: Reinstall the controller module

Reinstall the controller module into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Turn the controller module over and align the end with the opening in the chassis.
4. Gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

5. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

6. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller begins to boot as soon as it is completely seated in the chassis.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.

Step 5: Restore automatic giveback and AutoSupport

Restore automatic giveback and AutoSupport if they have been disabled.

1. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
2. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - FAS2820

To replace the chassis, you must move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving all drives and controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-controller cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - FAS2820

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace
   chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - FAS2820

Move the power supplies, hard drives, and controller module or modules from the impaired chassis to the replacement chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the replacement chassis of the same model as the impaired chassis.

Step 1: Move a power supply

Moving out a power supply when replacing a chassis involves turning off, disconnecting, and removing the power supply from the impaired chassis and installing and connecting it on the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.
4. Use the cam handle to slide the power supply out of the system.

   

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

   

   Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

Step 2: Remove the controller module

Remove the controller module or modules from the impaired chassis.

1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

2. Remove and set aside the cable management devices from the left and right sides of the controller module.
3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

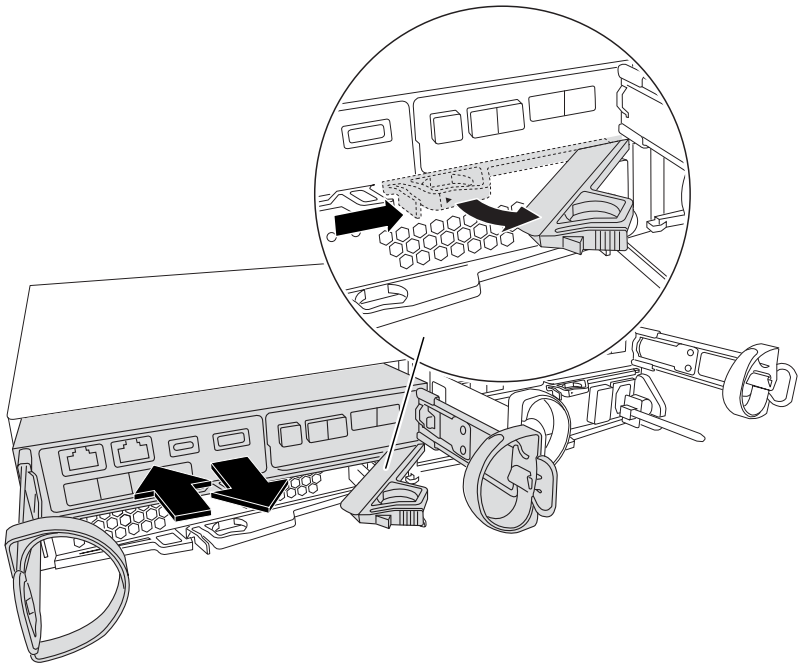

4. Set the controller module aside in a safe place.
5. Repeat these steps for the second controller module in the chassis.

Step 3: Move drives to the replacement chassis

Move the drives from each drive bay opening in the impaired chassis to the same bay opening in the replacement chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button on the opposite side of the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the impaired chassis with the same bay opening in the replacement chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate to the closed position.

5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It click when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 4: Replace a chassis from within the equipment rack or system cabinet

Remove the existing chassis from the equipment rack or system cabinet and install the replacement chassis in the equipment rack or system cabinet.

1. Remove the screws from the chassis mount points.
2. With the help of two or three people, slide the impaired chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller

Install the controller module and any other components into the replacement chassis, boot it tto Maintenance mode.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Repeat the preceding steps for the second controller in the replacement chassis.
4. Complete the installation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.

d. Repeat the preceding steps for the second controller module in the replacement chassis.

5. Connect the power supplies to different power sources, and then turn them on.
6. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

      If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

   b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - FAS2820

Verify the HA state of the chassis bring up the system,and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis based on the system's existing configuration: `ha-config modify chassis` *`ha-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. Exit Maintenance mode: `halt`. The LOADER prompt appears.
5. Boot the controller modules.

Step 2: Bring up the system

1. If you have not done so, plug the power cables back into the PSUs.
2. Turn on the PSUs by toggling the rocker switched to **ON**, and wait for the controllers to power up completely.
3. Check the front and the back of the chassis and controllers for any fault lights after power up.
4. Connect to the SP or BMC IP address of the nodes via SSH. This will be the same address used to shut

down the nodes.

5. Perform additional health checks as described in How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP
6. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

As a best practice, you should do the following:

- Resolve any Active IQ Wellness Alerts and Risks (Active IQ will take time to process post-power up AutoSupports - expect a delay in results)
- Run Active IQ Config Advisor
- Check system health using How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - FAS2820

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- This procedure includes steps for automatically or manually reassigning drives to the *replacement* controller, depending on your system’s configuration.

  You should perform the drive reassignment as directed in the procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - FAS2820

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the controller module hardware - FAS2820

Replace the impaired controller module hardware by removing the impaired controller, moving FRU components to the replacement controller module, installing the replacement controller module in the chassis, and then booting the replacement controller module.

Animation - Replace a controller module

Step 1: Remove controller module

Remove the impaired controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.
4. If you left the SFP modules in the system after removing the cables, move them to the replacement controller module.
5. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

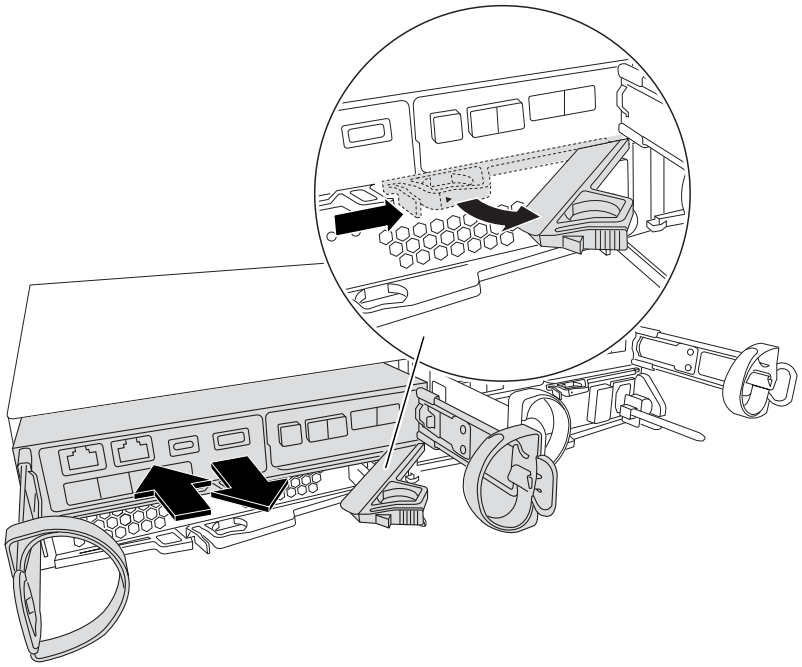

6. Turn the controller module over and place it on a flat, stable surface.
7. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.

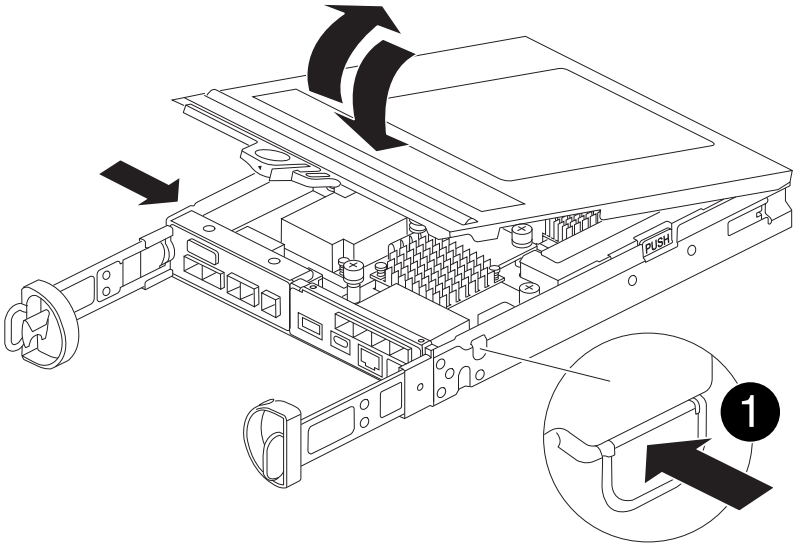

1	Controller module cover release button

Step 2: Move the NVMEM battery

Remove the NVMEM battery from the impaired controller module and install it into the replacement controller module.

Do not plug the NVMEM battery in until directed to do so.

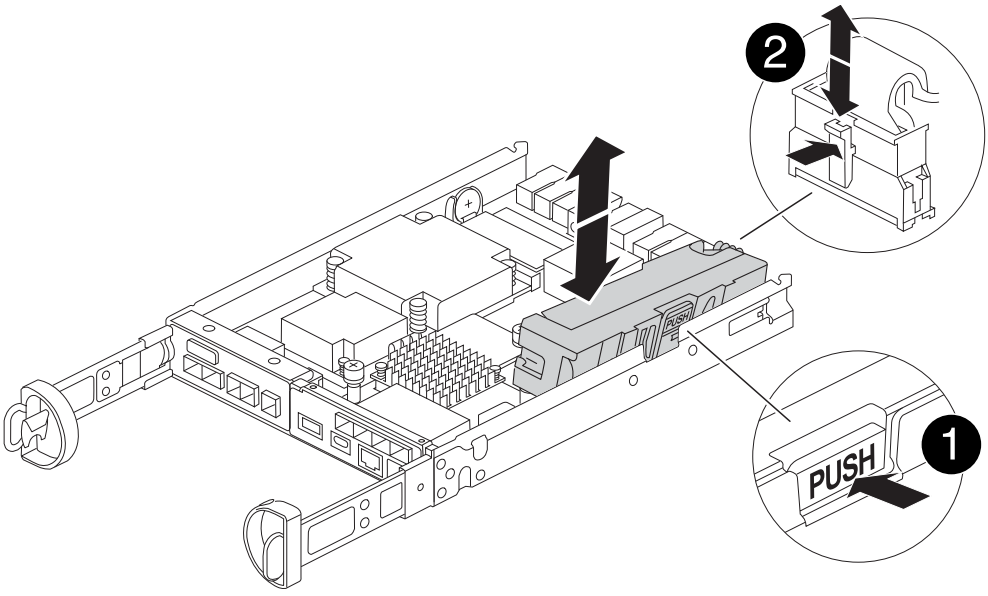

1	NVMEM battery release button
2	NVMEM battery plug

1. Remove the battery from the controller module:

a. Press the blue button on the side of the controller module.

b. Slide the battery up until it clears the holding brackets, and then lift the battery out of the controller module.

c. Unplug the battery plug by squeezing the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.

2. Move the battery to the replacement controller module and install it:

a. Aligning the battery with the holding brackets on the sheet metal side wall.

b. Slide the battery pack down until the battery latch engages and clicks into the opening on the side wall.

> Do not plug the battery in yet. You will plug it in once the rest of the components are moved to the replacement controller module.

Step 3: Remove the mezzanine card

Remove the IO Plate and PCIe mezzanine card from the impaired controller module.

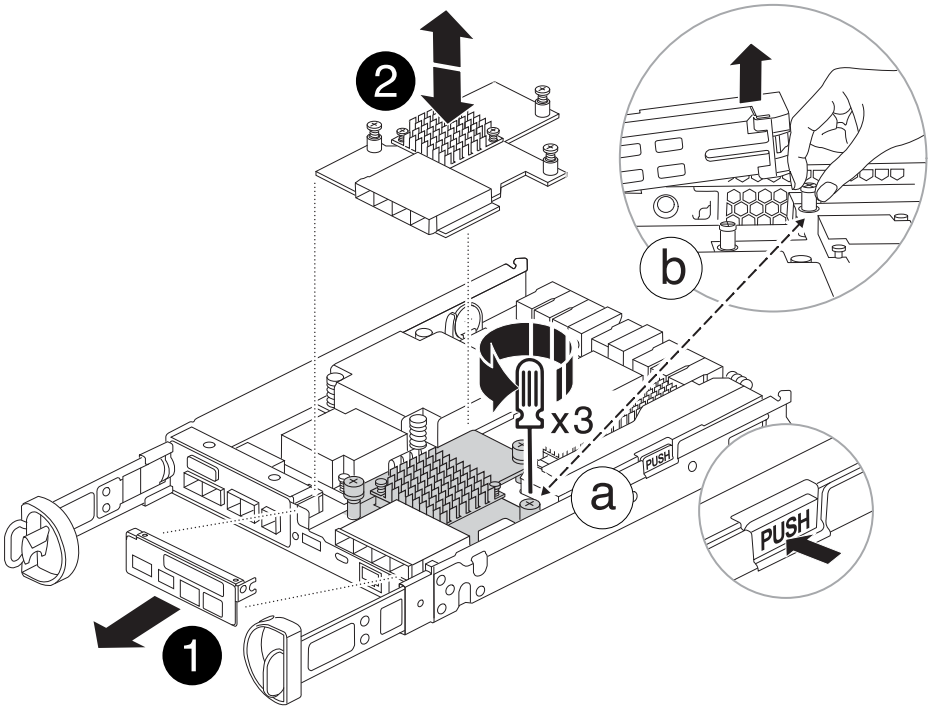

1	IO Plate
2	PCIe mezzanine card

1. Remove the IO Plate by sliding it straight out from the controller module.
2. Loosen the thumbscrews on the mezzanine card.

> You can loosen the thumbscrews with your fingers or a screwdriver.

3. Lift the mezzanine card straight up and set it aside on an anti-static surface.

Step 4: Move the boot media

Remove the boot media from the impaired controller module and install it in the replacement controller module.

1. After removing the mezzanine card, locate the boot media using the following illustration or the FRU map on the controller module:

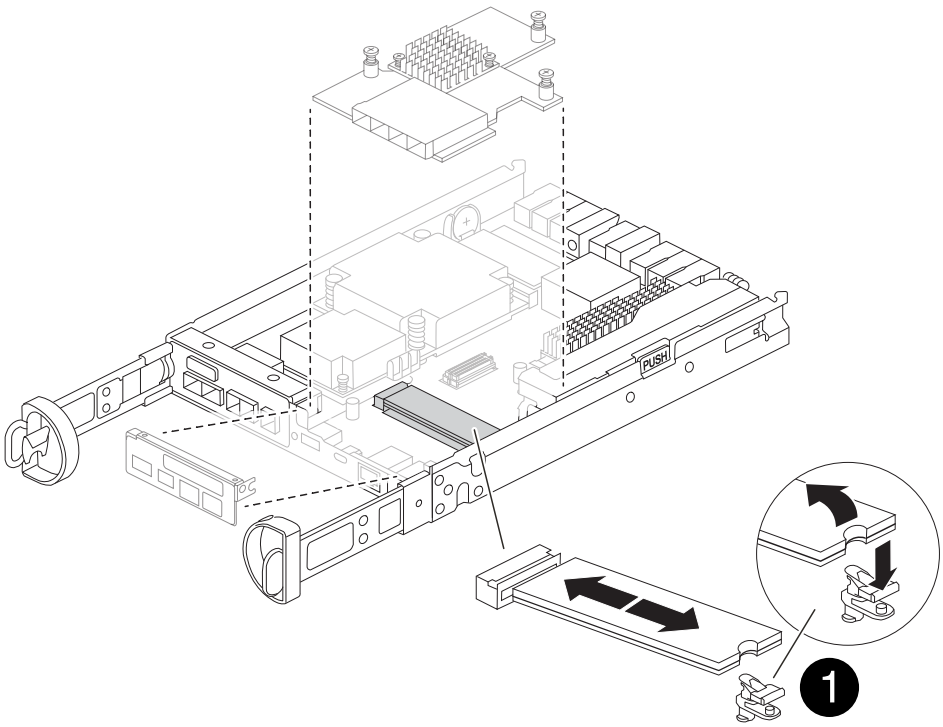

1	Boot media release button

2. Remove the boot media:
   a. Press the blue button on the boot media housing to release the boot media from its housing.
   b. Rotate the boot media up, and then gently pull it straight out of the boot media socket.

   > Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Install the the boot media to the replacement controller module:
   a. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
   b. Check the boot media to make sure that it is seated squarely and completely in the socket.

      If necessary, remove the boot media and reseat it into the socket.

   c. Push the blue locking button on the boot media housing, rotate the boot media all the way down, and then release the locking button to lock the boot media in place.

Step 5: Install the mezzanine card in the replacement controller

Install the mezzanine card in the replacement controller module.

1. Reinstall the mezzanine card:
   a. Align mezzanine card with the socket on the motherboard.
   b. Gently push down on the card to seat the card in the socket.
   c. Tighten the three thumbscrews on the mezzanine card.
2. Reinstall the IO Plate.

Step 6: Move the DIMMs

Remove the DIMMs from the impaired controller module and install them into the replacement controller module.

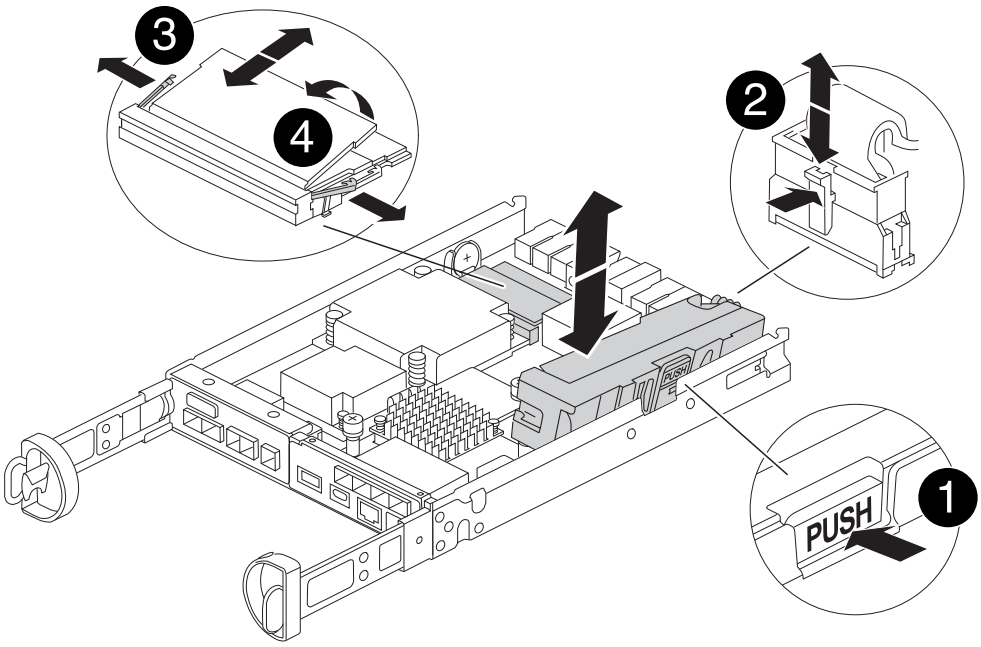

1	DIMM locking latches
2	DIMM

1. Locate the DIMMs on your controller module

   Note the location of the DIMM in the sockets so that you can insert the DIMM in the same location in the replacement controller module and in the proper orientation.

2. Remove the DIMMs from the impaired controller module:
   a. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM.

      The DIMM will rotate up a little.

   b. Rotate the DIMM as far as it will go, and then slide the DIMM out of the socket.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

3. Verify that the NVMEM battery is not plugged into the replacement controller module.
4. Install the DIMMs in the replacement controller in the same place they were in the impaired controller:
   a. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.

      The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

      Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.
5. Repeat these steps for the other DIMM.

Step 7: Move a caching module

Remove the caching module from the impaired controller module install it into replacement controller module.

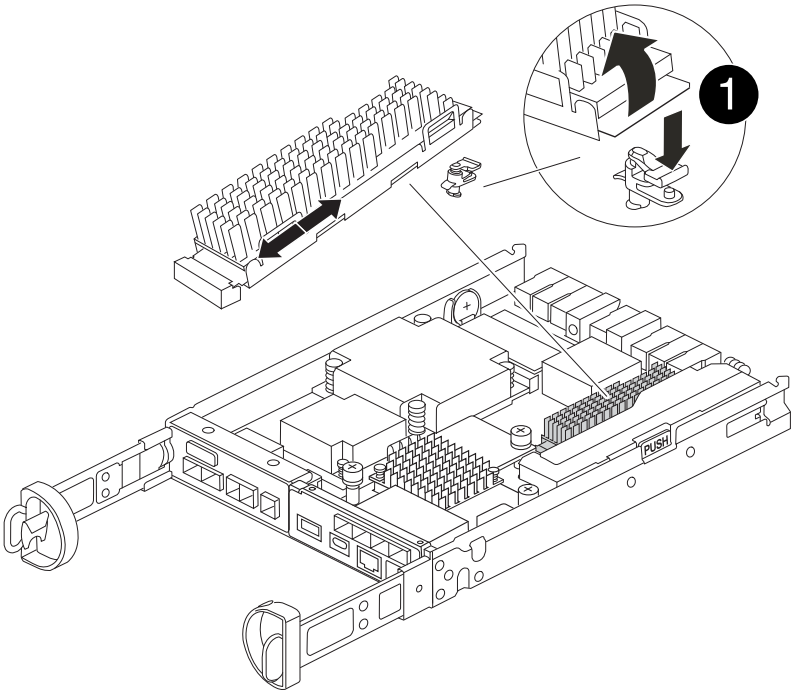

1	Caching module locking button

1. Locate the caching module near the rear of the controller module and remove it:
   a. Press the blue locking button and rotate the caching module upward.
   b. Gently pull the caching module straight out of the housing.
2. Install the caching module in the replacement controller module:
   a. Align the edges of the caching module with the socket in the housing, and then gently push it into the socket.

b. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

c. Push the blue locking button, rotate the caching module all the way down, and then release the locking button to lock the caching module in place.

3. Plug in the NVMEM battery.

   Make sure that the plug locks down into the battery power socket on the motherboard.

   If plugging in the battery is difficult, remove the battery from the controller module, plug it in, and then reinstall the battery into the controller module.

4. Reinstall the controller module cover.

Step 8: Install the NV battery

Install the NV battery into the replacement controller module.

1. Plug the battery plug back into the socket on the controller module.

   Make sure that the plug locks down into the battery socket on the motherboard.

2. Aligning the battery with the holding brackets on the sheet metal side wall.
3. Slide the battery pack down until the battery latch engages and clicks into the opening on the side wall.
4. Reinstall the controller module cover and lock it into place.

Step 9: Install the controller

Install the replacement controller module into the system chassis and boot ONTAP.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Turn the controller module.
4. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

The controller begins to boot as soon as it is seated in the chassis.

b. If you have not already done so, reinstall the cable management device.

c. Bind the cables to the cable management device with the hook and loop strap.

You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the `update_flash` command, and then enter `bye -g` to reboot the system.

Important: During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID. Respond `y` to this prompt.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down. Respond `y` to this prompt.

Restore and verify the system configuration - FAS2820

After completing the hardware replacement and booting the replacement controller, verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the controller does not match your system configuration, set the HA state for the replacement controller module: `ha-config modify controller HA-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`

     a. Confirm that the setting has changed: `ha-config show`

3. Reboot the controller module.

   During the boot process, you might see the following prompts:

   - A prompt warning of a system ID mismatch and asking to override the system ID.
   - A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
     You can safely respond `y` to these prompts.

Recable the system and reassign disks - FAS2820

To complete the replacement procedure and restore your system to full operation, you must recable the storage, confirm disk reassignment, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`
   c. Wait for the `savecore` command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system`

```
node run -node local-node-name partner savecore -s
```

d. Return to the admin privilege level: `set -privilege admin`

5. Give back the controller:

a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

The *replacement* controller takes back its storage and completes booting.

If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

> If the giveback is vetoed, resolve the veto issue. If the veto is not critical to resolve, you can override the veto.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

6. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Complete system restoration - FAS2820

Restore your system to full operation by restoring the NetApp Storage Encryption or Volume Encryption configurations (if necessary), and installing licenses for the replacement controller, and returning the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS2820

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Animation - Replace a DIMM

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove controller module

Remove the controller module from the system and then remove the controller module cover.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
3. Remove and set aside the cable management devices from the left and right sides of the controller module.
4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

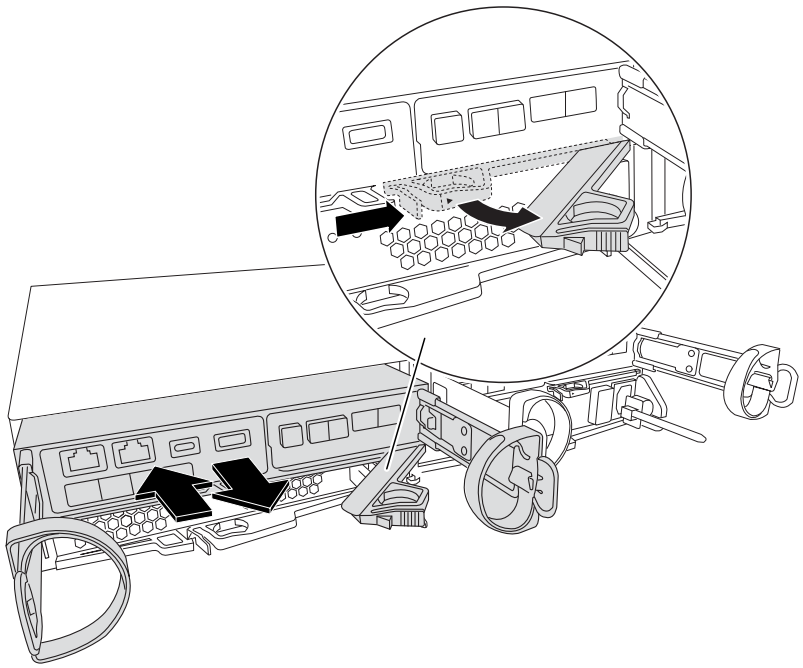

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.

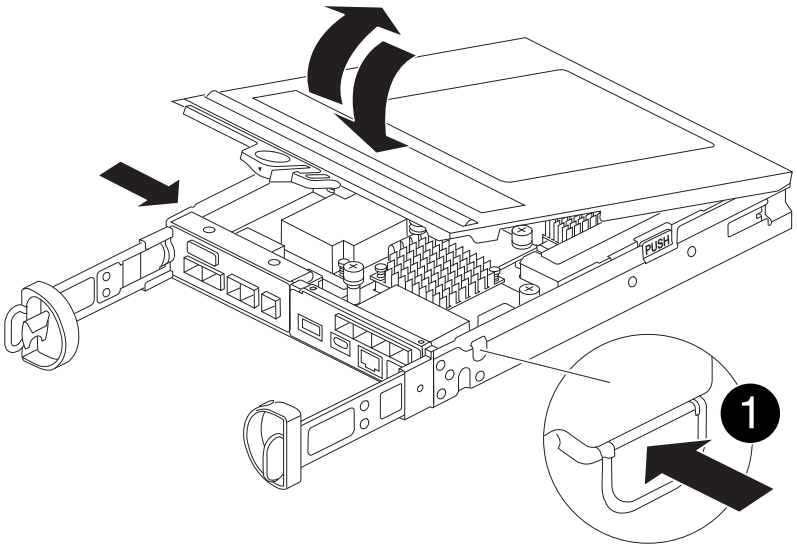

1	Controller module cover release button

Step 3: Replace the DIMMs

Locate the DIMM inside the controller, remove it, and replace it.

Before replacing a DIMM, you need to unplug the NVMEM battery from the controller module.

Steps

1. If you are not already grounded, properly ground yourself.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

   

2. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
3. If the NVMEM LED is flashing, there is data in the NVMEM and you must disconnect the battery to clear the memory:
   a. Remove the battery from the controller module by pressing the blue button on the side of the controller module.
   b. Slide the battery up until it clears the holding brackets, and then lift the battery out of the controller module.
   c. Locate the battery cable, press the clip on the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

d. Confirm that the NVMEM LED is no longer lit.

e. Reconnect the battery connector and recheck the LED on the back of the controller.

f. Unplug the battery cable.

4. Locate the DIMMs on your controller module.
5. Note the orientation and location of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
6. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   The DIMM will rotate up a little.

7. Rotate the DIMM as far as it will go, and then slide the DIMM out of the socket.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

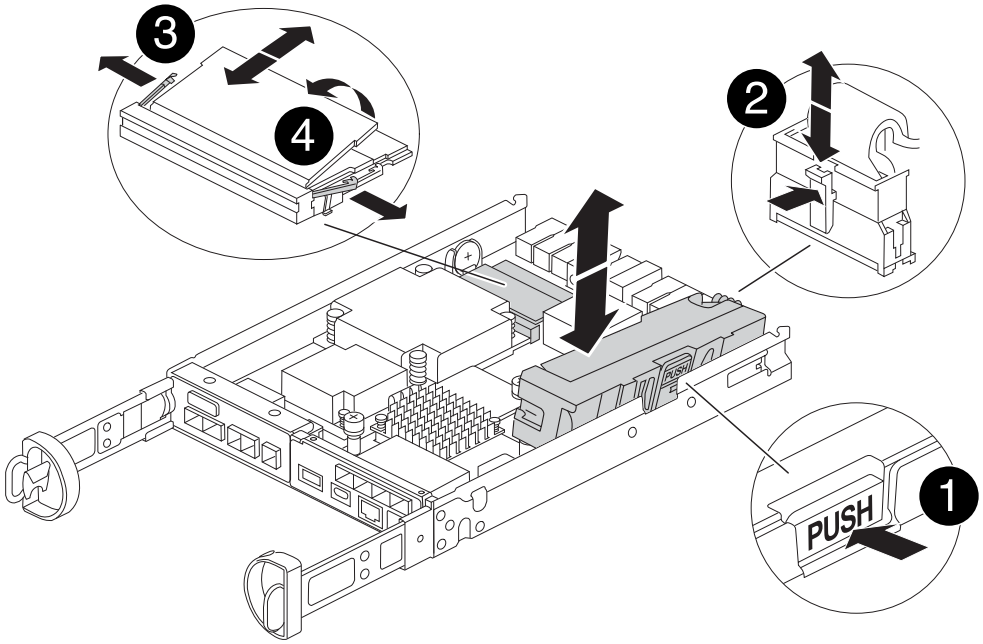

1	NVRAM battery release button
2	NVRAM battery plug
3	DIMM ejector tabs
4	DIMMs

8. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

9. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

10. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
11. Reconnect the NVMRM battery:
    a. Plug in the NVRAM battery.

       Make sure that the plug locks down into the battery power socket on the motherboard.

    b. Align the battery with the holding brackets on the sheet metal side wall.
    c. Slide the battery pack down until the battery latch engages and clicks into the opening on the side wall.
12. Reinstall the controller module cover.

Step 4: Reinstall the controller module

Reinstall the controller module into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Turn the controller module over and align the end with the opening in the chassis.
4. Gently push the controller module halfway into the system.Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

5. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

6. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller begins to boot as soon as it is seated in the chassis.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
7. Reboot the controller module.

During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
  You can safely respond `y` to these prompts.

Step 5: Restore automatic giveback and AutSupport

Restore automatic giveback and AutoSupport if they have been disabled.

1. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
2. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - FAS2820

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.

- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   > The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace the NVMEM battery - FAS2820

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact NetApp Support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove and open the controller module

Remove and open the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.
4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module half-way out of the chassis.

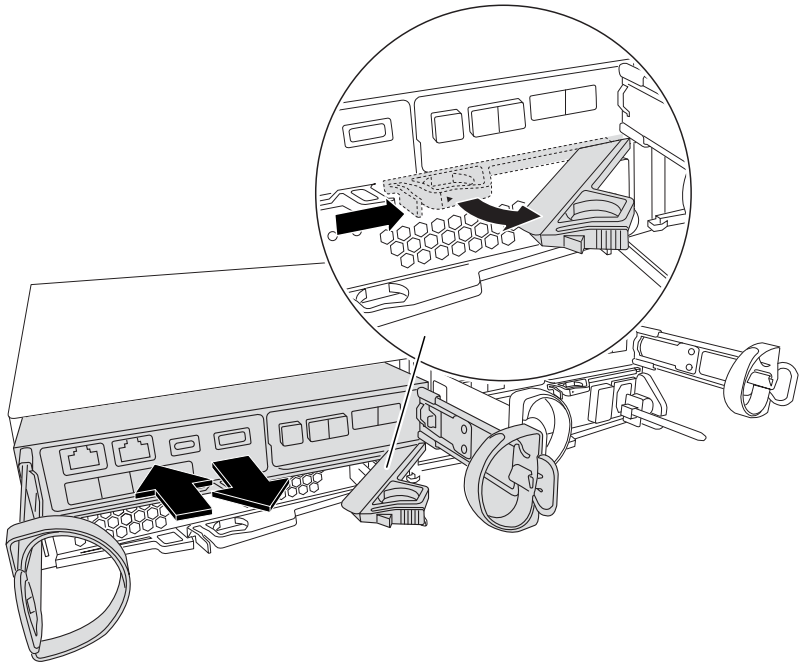

5. Check the NVMEM LED located on the back of the controller module. Look for the NV icon:

The green NV LED on the faceplate will start flashing when power is removed from the controller if the system was in the "waiting for giveback" state, or the system was not taken over or halted properly (uncommitted data). If the impaired controller module was not successfully taken over by the partner controller module, contact NetApp Support

   - If the green NV status LED begins flashing when the controller module is removed from the chassis:
     - Confirm that the controller had a clean takeover by the partner controller module or the impaired controller shows *waiting for giveback*, the flashing LED can be ignored and you can complete removing the impaired controller from the chassis.

◦ If the green NV LED if off, you can complete removing the impaired controller from the chassis.

Step 3: Replace the NVMEM battery

Remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the controlle module from the chassis.
3. Turn the controller module over and place it on a flat, stable surface.
4. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.

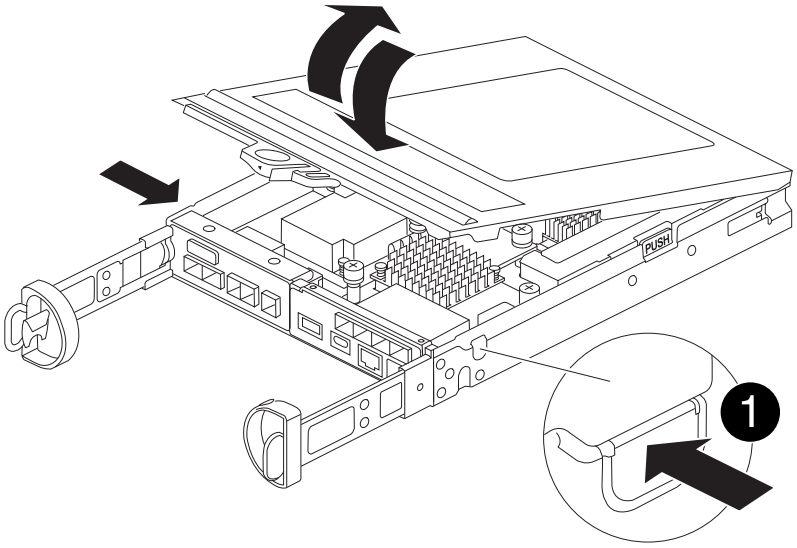

5. Locate the NVMEM battery in the controller module.

   Animation - Replace the NV battery

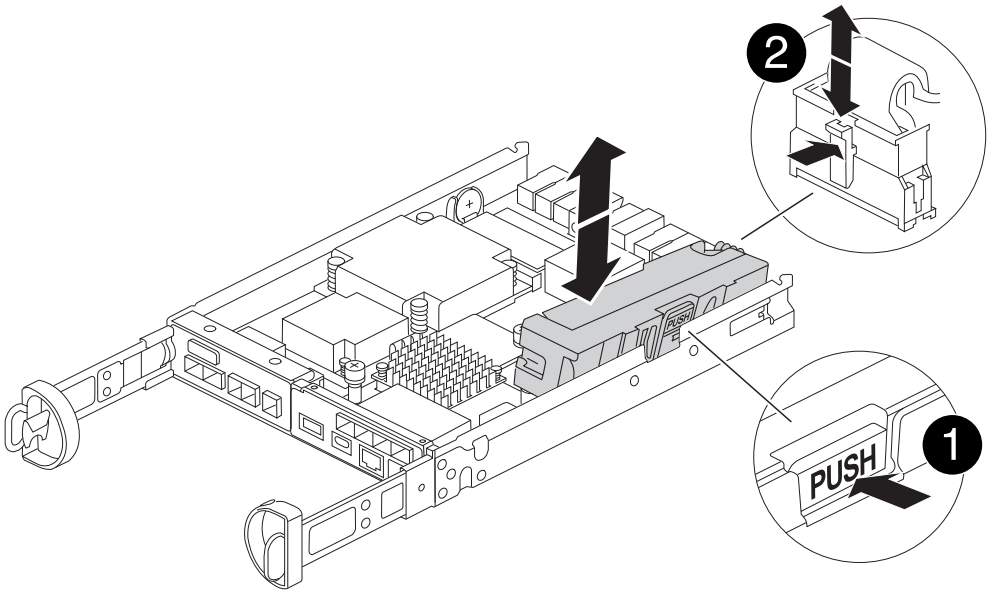

1	Battery release tab
2	Battery power connector

6. Remove the failed battery from the controller module:
   a. Press the blue button on the side of the controller module.
   b. Slide the battery up until it clears the holding brackets, and then lift the battery out of the controller module.
   c. Unplug the battery from the controller module
7. Remove the replacement battery from its package.
   Install the replacement battery:
   a. Plug the battery plug back into the socket on the controller module.

      Make sure that the plug locks down into the battery socket on the motherboard.

   b. Aligning the battery with the holding brackets on the sheet metal side wall.
   c. Slide the battery pack down until the battery latch engages and clicks into the opening on the side wall.
8. Reinstall the controller module cover and lock it into place.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Turn the controller module over and align the end with the opening in the chassis.
4. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

5. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

6. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller begins to boot as soon as it is seated in the chassis.

b. If you have not already done so, reinstall the cable management device.

c. Bind the cables to the cable management device with the hook and loop strap.

7. Reboot the controller module.

   During the boot process, you might see the following prompts:

   - A prompt warning of a system ID mismatch and asking to override the system ID.
   - A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
     You can safely respond `y` to these prompts.

Step 5: Restore automatic giveback and AutSupport

Restore automatic giveback and AutoSupport if they have been disabled.

1. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
2. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a mezzanine card - FAS2820

Replace the mezzanine card by disconnecting the cables and any SFP and QSFP modules from the card, replace the failed mezzanine card, and then recable the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Animation - Replace the mezzanine card

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove the controller module

Remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
3. Remove and set aside the cable management devices from the left and right sides of the controller module.
4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

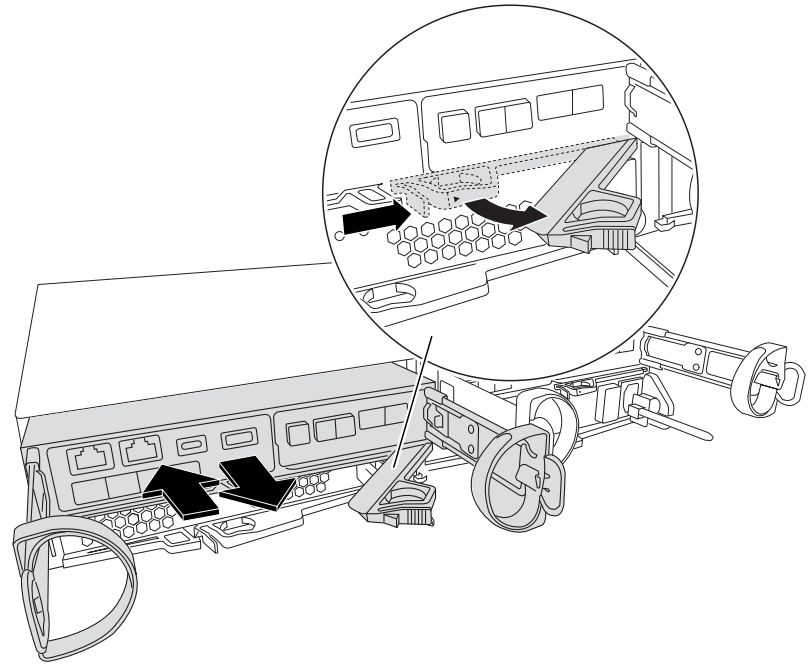

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.

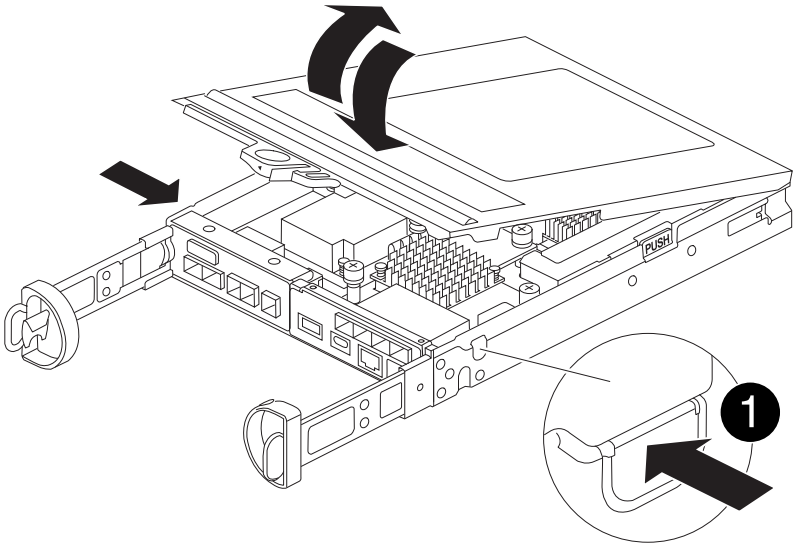

1	Controller module cover release button

Step 3: Replace the mezzanine card

Replace the mezzanine card.

1. If you are not already grounded, properly ground yourself.
2. Remove the mezzanine card using the following illustration or the FRU map on the controller module:

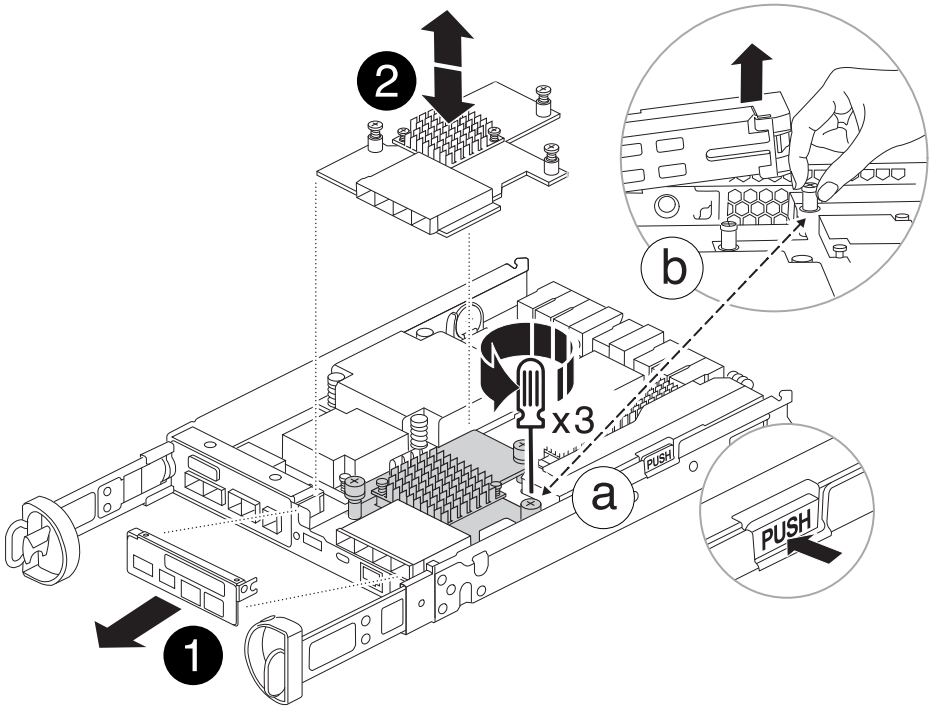

1	IO Plate
2	PCIe mezzanine card

a. Remove the IO Plate by sliding it straight out from thr controller module.

b. Loosen the thumbscrews on the mezzanine card and lift the mezzanine card straight up.

> You can loosen the thumbscrews with your fingers or a screwdriver. If you use your fingers, you might need to rotate the NV battery up for better finger purchase on the thumbscrew next to it.

3. Reinstall the mezzanine card:

   a. Align the socket on the replacement mezzanine card plug with the socket on the motherboard, and then gently seat the card squarely into the socket.

   b. Tighten the three thumbscrews on the mezzanine card.

   c. Reinstall the IO Plate.

4. Reinstall the controller module cover and lock it into place.

Step 4: Install the controller module

Reinstall the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Turn the controller module over and align the end with the opening in the chassis.

4. Gently push the controller module halfway into the system.Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

5. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

6. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller begins to boot as soon as it is seated in the chassis.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
9. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - FAS2820

Swapping out a power supply involves turning off, disconnecting, and removing the impaired power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

  It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- Power supplies are auto-ranging.

Animation - Replace power supply

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.

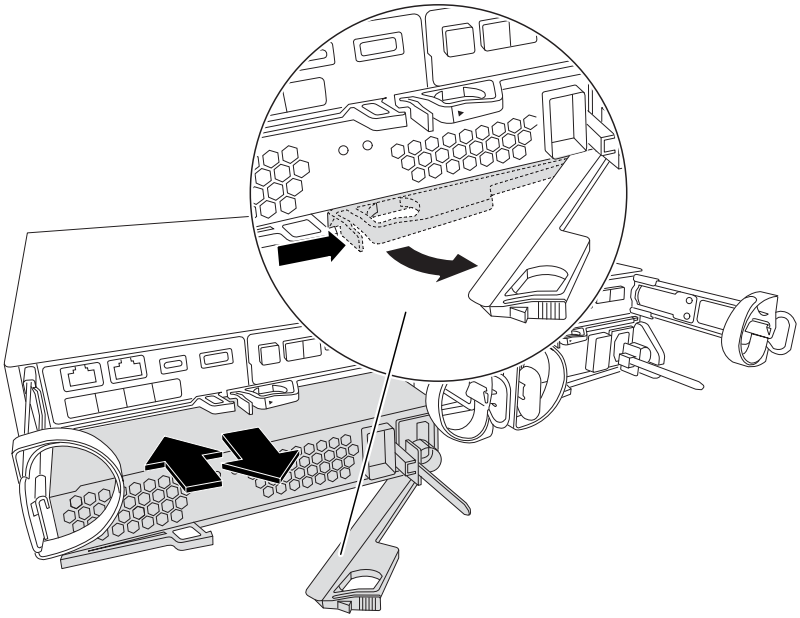

5. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.
7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
9. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply and the power source.

b. Secure the power cable to the power supply using the power cable retainer.

Once power is restored to the power supply, the status LED should be green.

10. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

    The power supply LEDs are lit when the power supply comes online.

11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - FAS2820

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove controller module

Remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.
4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

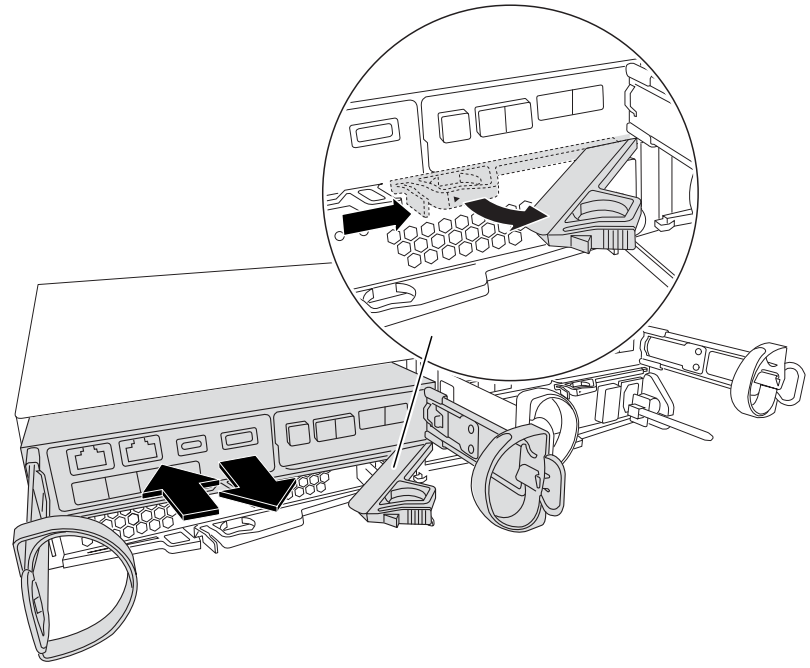

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.

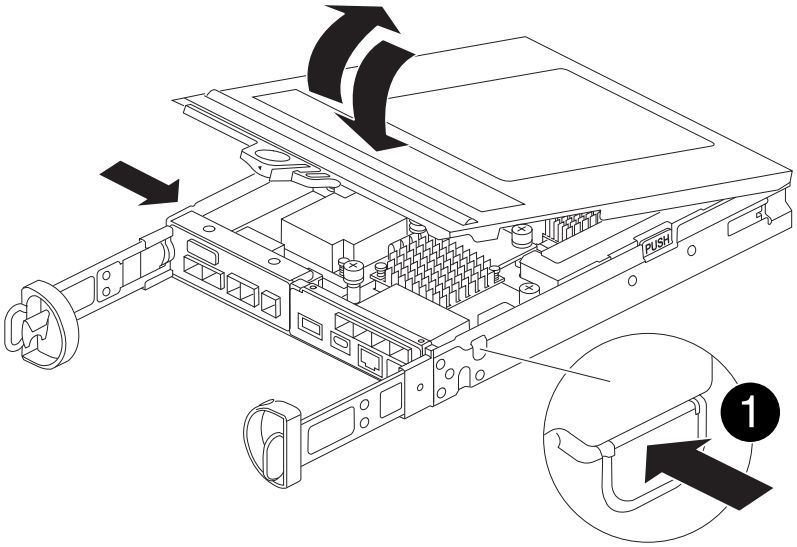

Step 3: Replace the RTC battery

Replace the RTC battery by locating it inside the controller and follow the specific sequence of steps.

Animation - Replace the RTC battery

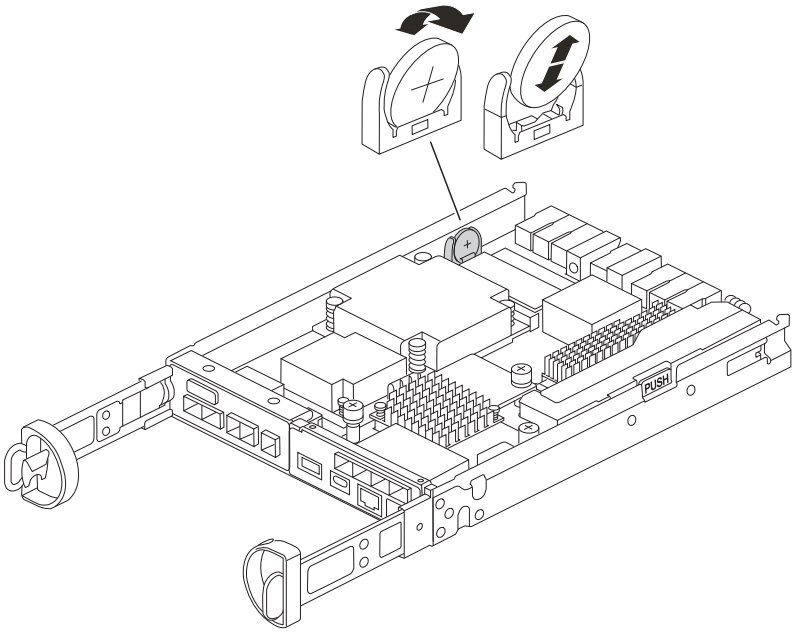

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.
3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
8. Reinstall the controller cover.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it to the LOADER prompt..

1. Turn the controller module over and align the end with the opening in the chassis.
2. Gently push the controller module halfway into the system.Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:

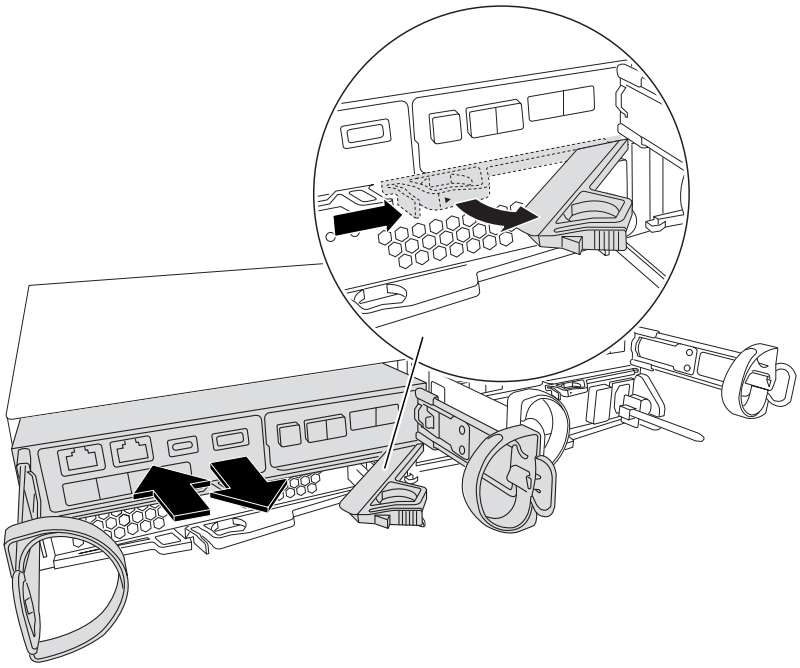

a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

> Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

The controller begins to boot as soon as it is seated in the chassis.

b. If you have not already done so, reinstall the cable management device.
c. Bind the cables to the cable management device with the hook and loop strap.
d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
e. Halt the controller at the LOADER prompt.

Step 5: Set time/date after RTC battery replacement

1. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
2. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
3. Return the controller to normal operation by giving back its storage: `storage failover giveback`

`-ofnode impaired_node_name`

4. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
5. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

FAS8300 and FAS8700 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

For MetroCluster configurations, see either:

- Install MetroCluster IP configuration
- Install MetroCluster Fabric-Attached configuration

Quick guide - FAS8300 and FAS8700

This guide gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Access the *Installation and Setup Instructions* PDF poster:

FAS8300 and FAS8700 Installation and Setup Instructions

Video steps - FAS8300 and FAS8700

The following video shows how to install and cable your new system.

Animation - FAS8300 and FAS8700 Install and setup instructions

Detailed guide - FAS8300 and FAS8700

This guide gives detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your system, you need to create an account, register the system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
100 GbE cable (QSF(28)	X66211A-05 (112-00595), 0.5m X66211A-1 (112-00573), 1m X66211A-2 (112-00574), 2m X66211A-5 (112-00574), 5m		Storage, cluster interconnect/HA, and Ethernet data (order-dependent)
25 GbE cable (SFP28s)	X66240-2 (112-00598), 2m X66240-5 (112-00639), 5m		GbE network connection (order-dependent)

Type of cable…	Part number and length	Connector type	For…
32 Gb FC (SFP+ Op)	X66250-2 (112-00342), 2m X66250-5 (112-00344), 5m X66250-15 (112-00346), 15m		FC network connection
Storage Cables	X66030A (112-00435), .5m X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		mini-SAS HD to mini-SAS HD cables (order-dependent)
Optical cables	X66250-2-N-C (112-00342)		16 Gb FC or 25GbE cables for mezzanine cards (order-dependent)
RJ-45 (order dependent)	X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network
Micro-USB console cable	Not applicable		Console connection used during software setup if laptop or console does not support network discovery.
Power cables	Not applicable		Powering up the system

4. Review the *NetApp ONTAP Confiuration Guide* and collect the required information listed in that guide.

   ONTAP Configuration Guide

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

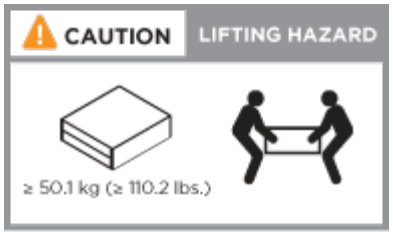

3. Attach cable management devices (as shown).

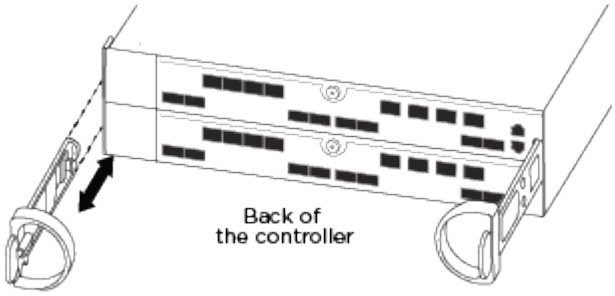

4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

If the port labels on the card are not visible, check the card installation orientation (the PCIe connector socket is on the left side of the card slot in the A400 and FAS8300/8700), and then look for the card,then look for the card, by part number, in the NetApp Hardware Universe for a graphic of the bezel which will show the port labels. The card part number can be found using the `sysconfig -a` command or on the system packing list.

Option 1: Cable a two-node switchless cluster

The optional data ports, optional NIC cards, and management ports on the controller modules are connected to switches. The cluster interconnect and HA ports are cabled on both controller modules.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.

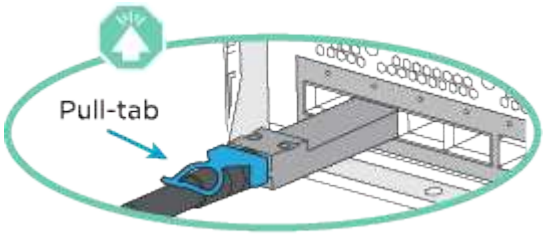

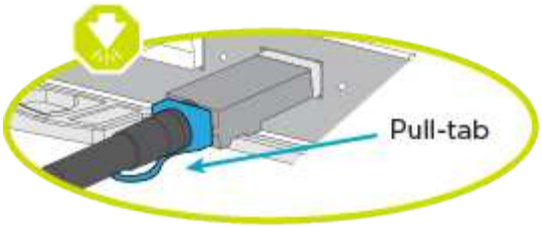

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Two-node switchless cluster cabling

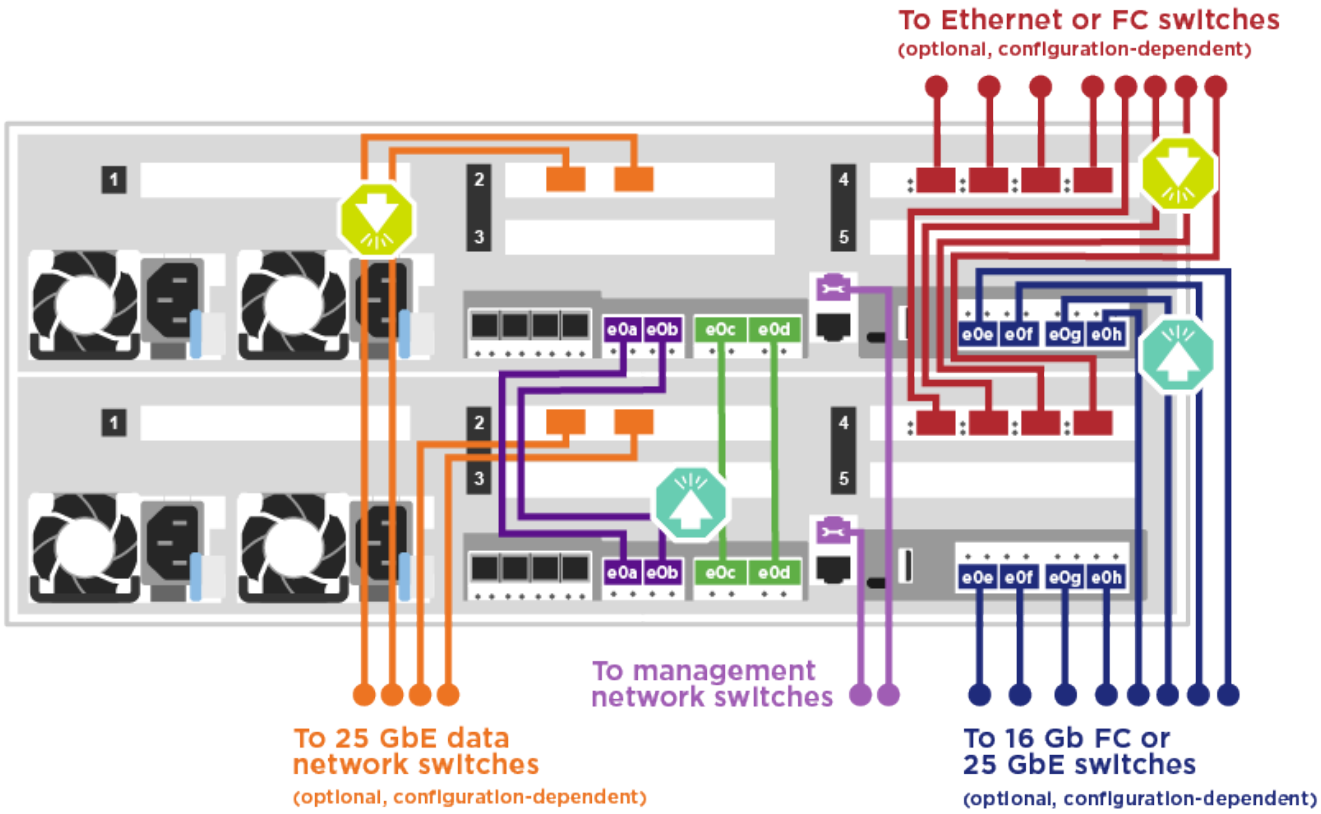

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Option 2: Cable a switched cluster

The optional data ports, optional NIC cards, mezzanine cards, and management ports on the controller modules are connected to switches. The cluster interconnect and HA ports are cabled on to the cluster/HA switch.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.

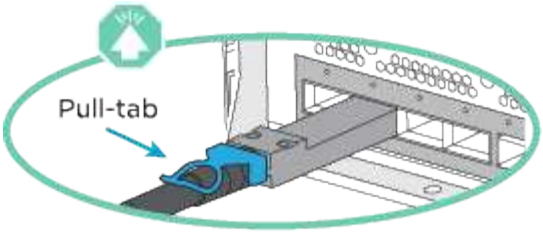

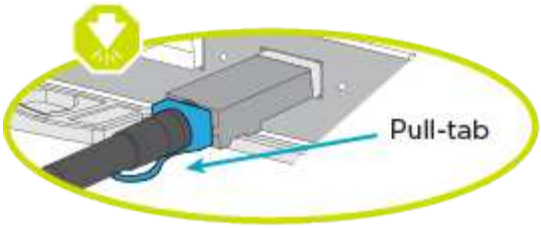

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Switched cluster cabling

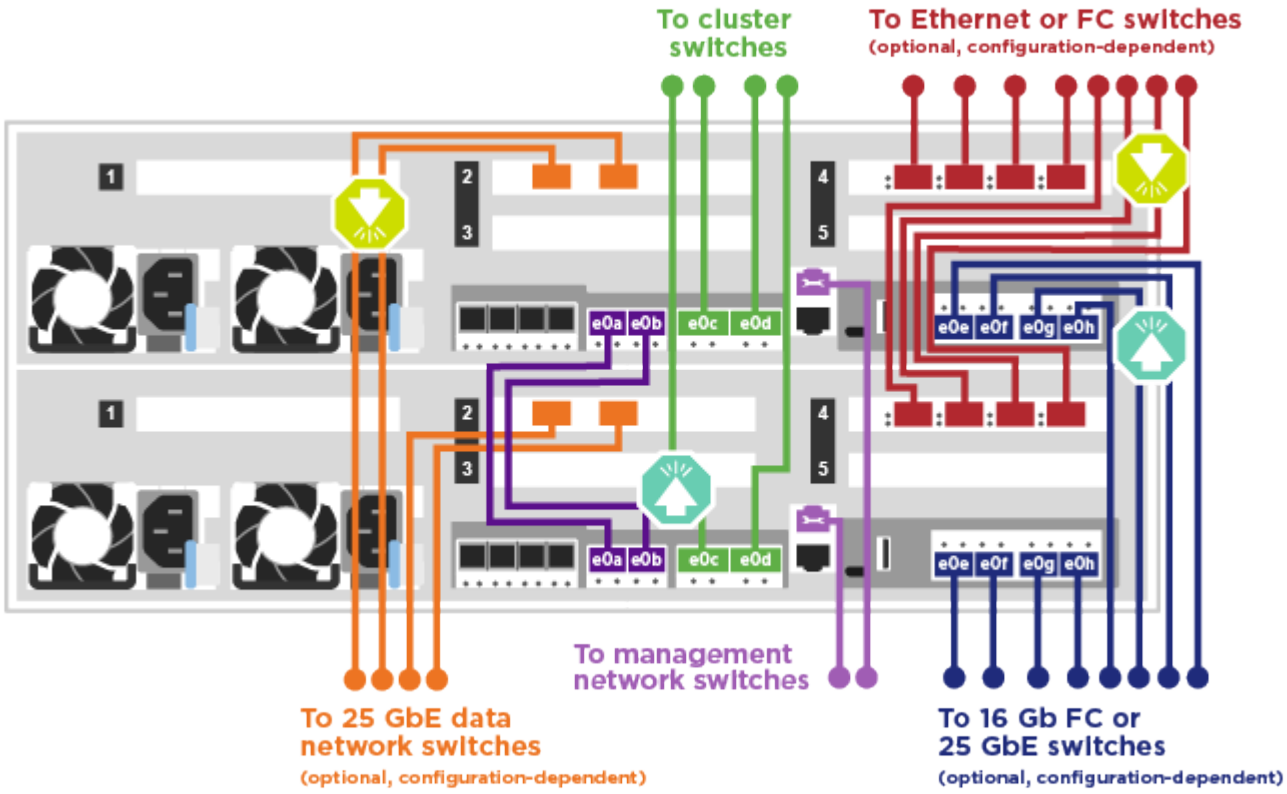

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Step 4: Cable controllers to drive shelves

Option 1: Cable the controllers to SAS drive shelves

You must cable each controller to the IOM modules on both SAS drive shelves.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the DS224-C are down.

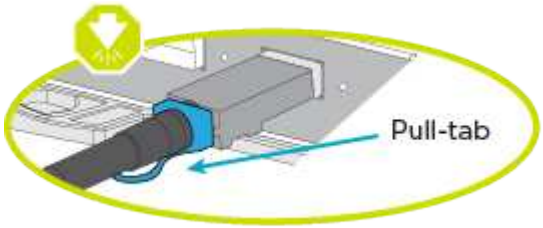

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following animation or illustration to cable your controllers to two drive shelves.

   Animation - Cable the controllers to SAS drive shelves

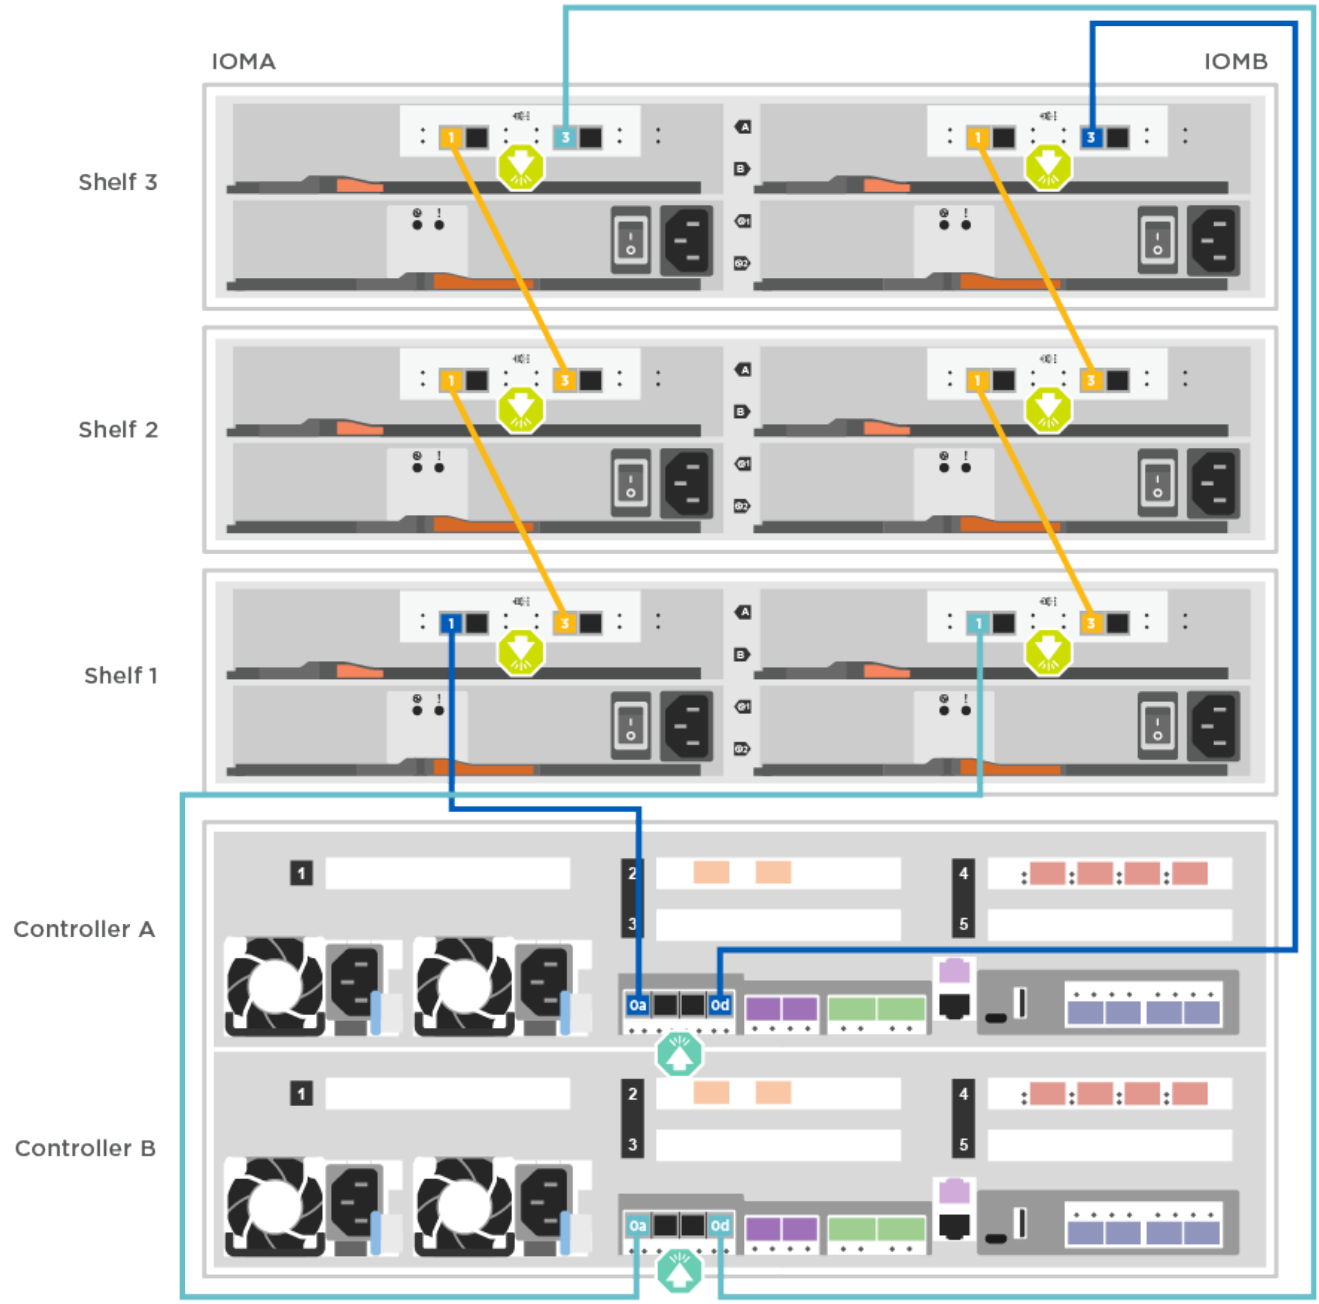

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Completing system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Use the following animation to set one or more drive shelf IDs:

Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

5. Select an ONTAP icon listed to discover:

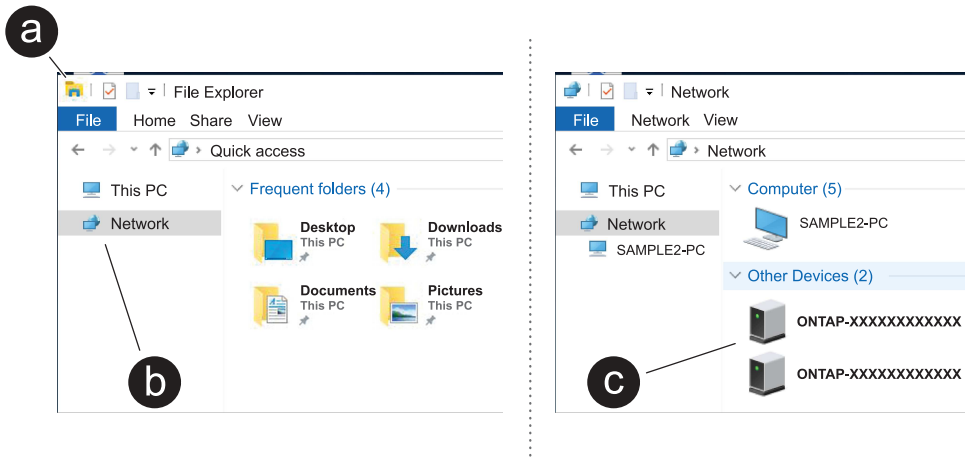

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

   System Manager opens.

6. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

   ONTAP Configuration Guide

7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet .
   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Use the following animation to set one or more drive shelf IDs:

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   FAS8300 and FAS8700 shown.

   Animation - Power on the controllers

   Initial booting may take up to eight minutes.

4. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

The format for the address is https://x.x.x.x.

b. Configure the system using the data you collected in the *NetApp ONTAP Configuration guide*.

ONTAP Configuration Guide

6. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain FAS8300 and FAS8700 hardware

Maintain the hardware of your FAS8300 and FAS8700 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS8300 and FAS8700 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS8300 and FAS8700 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the storage system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Caching module	You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.
NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.
PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - FAS8300 and FAS8700

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on the partner node to reinstall ONTAP on the replacement boot media in your FAS8300 or FAS8700 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1 **Review the boot media requirements**

Review the requirements for boot media replacement.

2 **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3 **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4 **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5 **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - FAS8300 and FAS8700

Before replacing the boot media in your FAS8300 or FAS8700, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.

- /cfcard/kmip/certs/CA.pem file.

- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - FAS8300 and FAS8700

Shut down the impaired controller in your FAS8300 or FAS8700 storage system to prevent data loss and ensure system stability when replacing the boot media.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - FAS8300 and FAS8700

The boot media in your FAS8300 or FAS8700 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

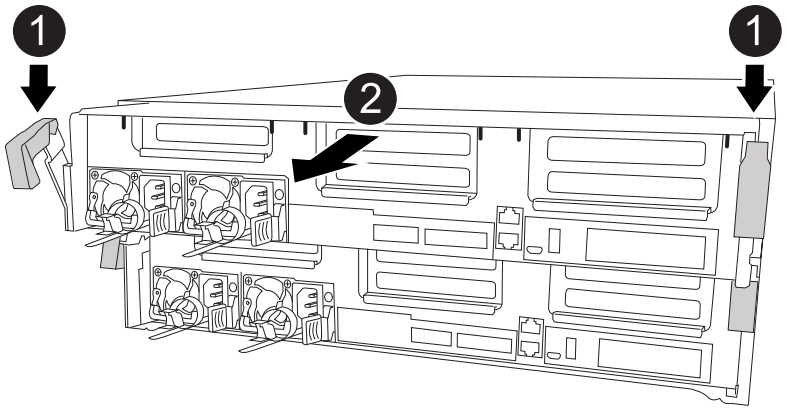

1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. Open the air duct:

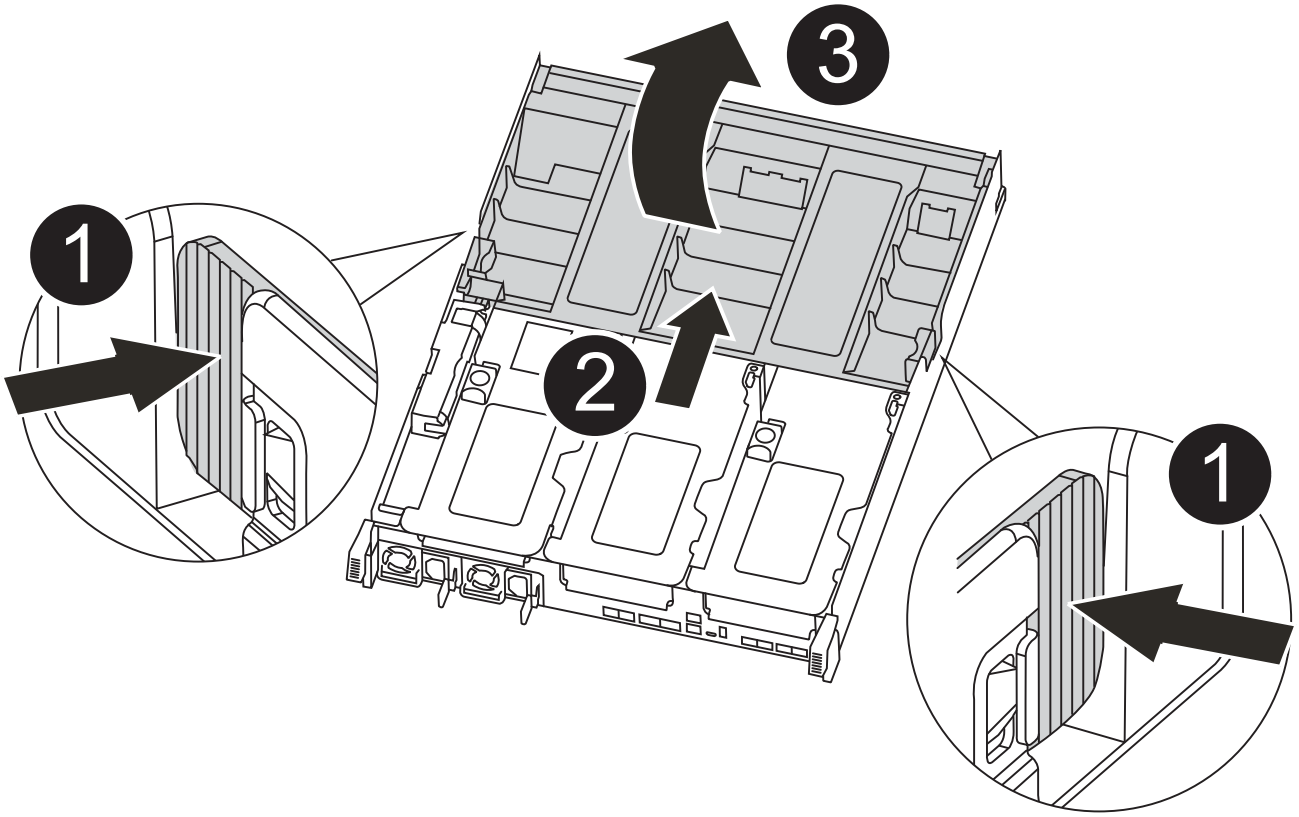

1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.

b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

9. Locate and remove the boot media from the controller module:

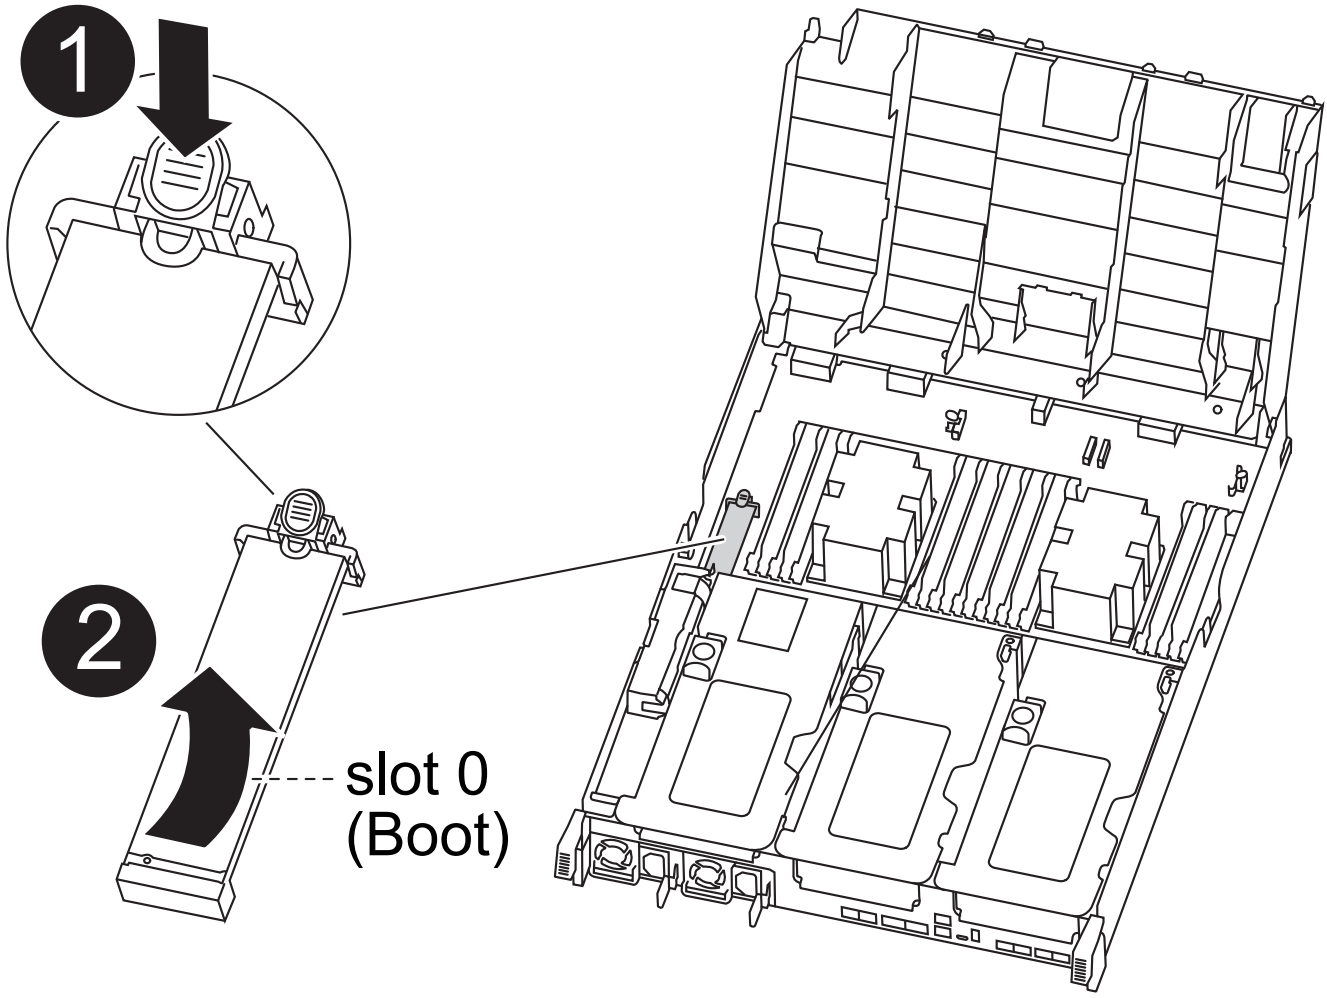

1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.

b. Rotate the boot media up and gently pull the boot media out of the socket.

10. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
11. Check the boot media to make sure that it is seated squarely and completely in the socket.

    If necessary, remove the boot media and reseat it into the socket.

12. Lock the boot media in place:
    a. Rotate the boot media down toward the motherboard.
    b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
    c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
13. Close the air duct.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - FAS8300 and FAS8700

After installing the new boot media device in your FAS8300 or FAS8700 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

> Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - FAS8300 and FAS8700

If a component in your FAS8300 or FAS8700 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - FAS8300 and FAS8700

Get started with replacing the boot media in your FAS8300 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - FAS8300 and FAS8700

Before replacing the boot media in your FAS8300 or FAS8700 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - FAS8300 and FAS8700

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - FAS8300 and FAS8700

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 3: Controller is in a two-node Metrocluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the boot media and prepare for manual boot recovery - FAS8300 and FAS8700

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

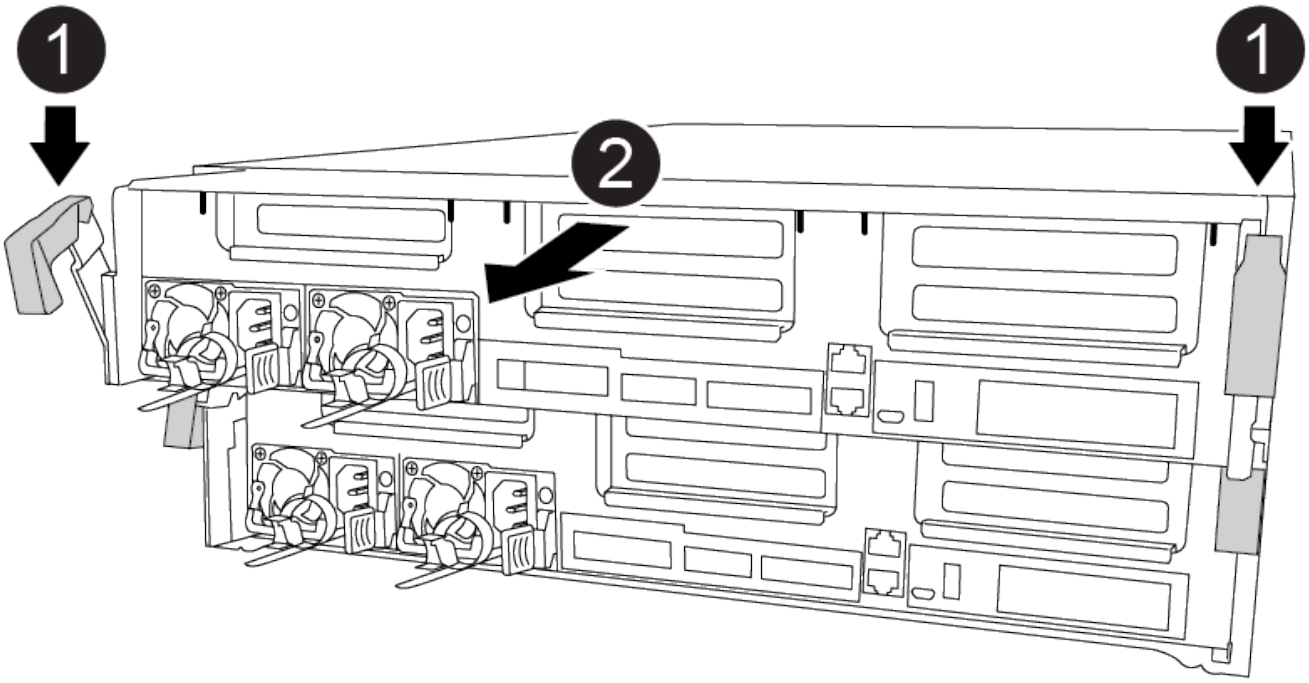

1	Locking latches
2	Slide controller out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 2: Replace the boot media

You must locate the boot media in the controller module (see the FRU map on the controller module), and then follow the directions to replace it.

Before you begin

Although the contents of the boot media is encrypted, it is a best practice to erase the contents of the boot media before replacing it. For more information, see the Statement of Volatility for your system on the NetApp Support Site.

> You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustrations, or the written steps to replace the boot media.

Animation - Replace the boot media

Steps

1. Open the air duct:

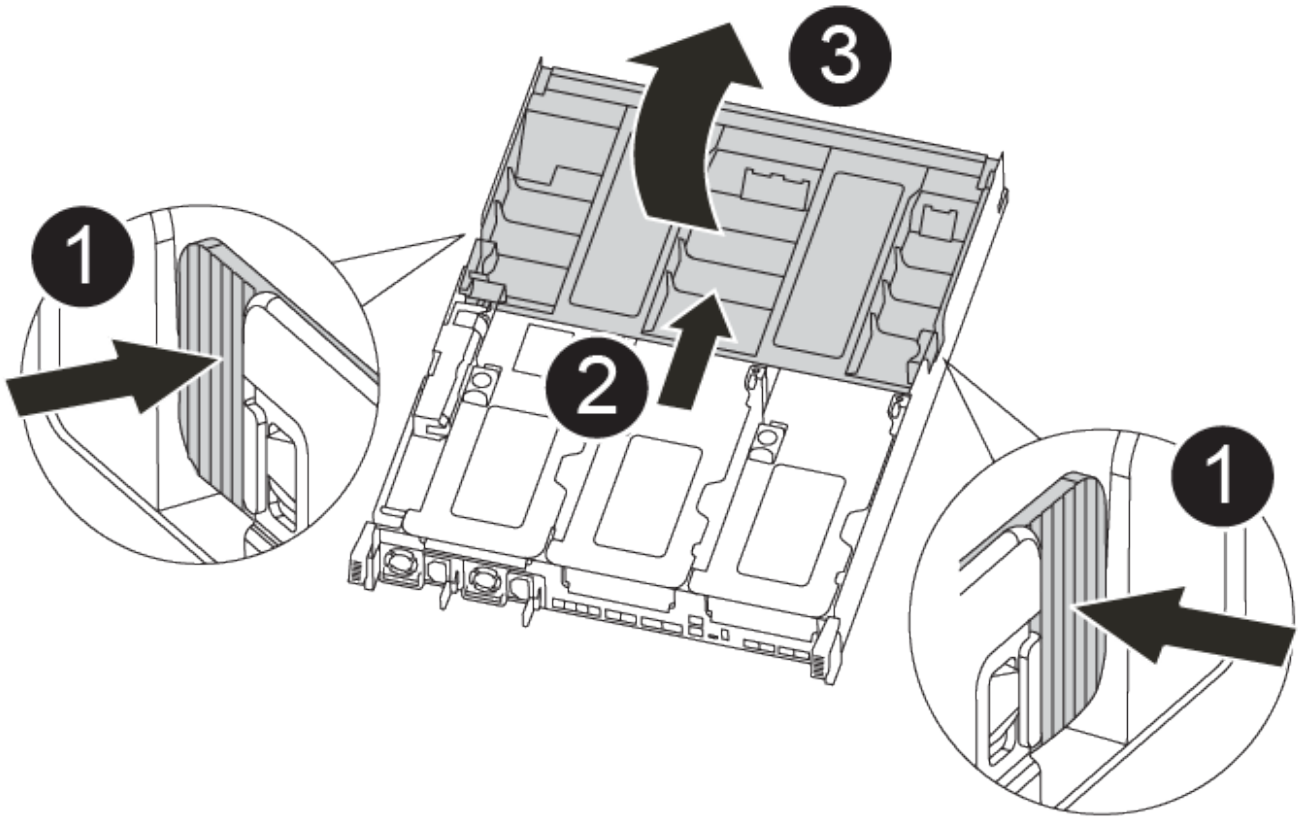

1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.

b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

2. Locate and remove the boot media from the controller module:

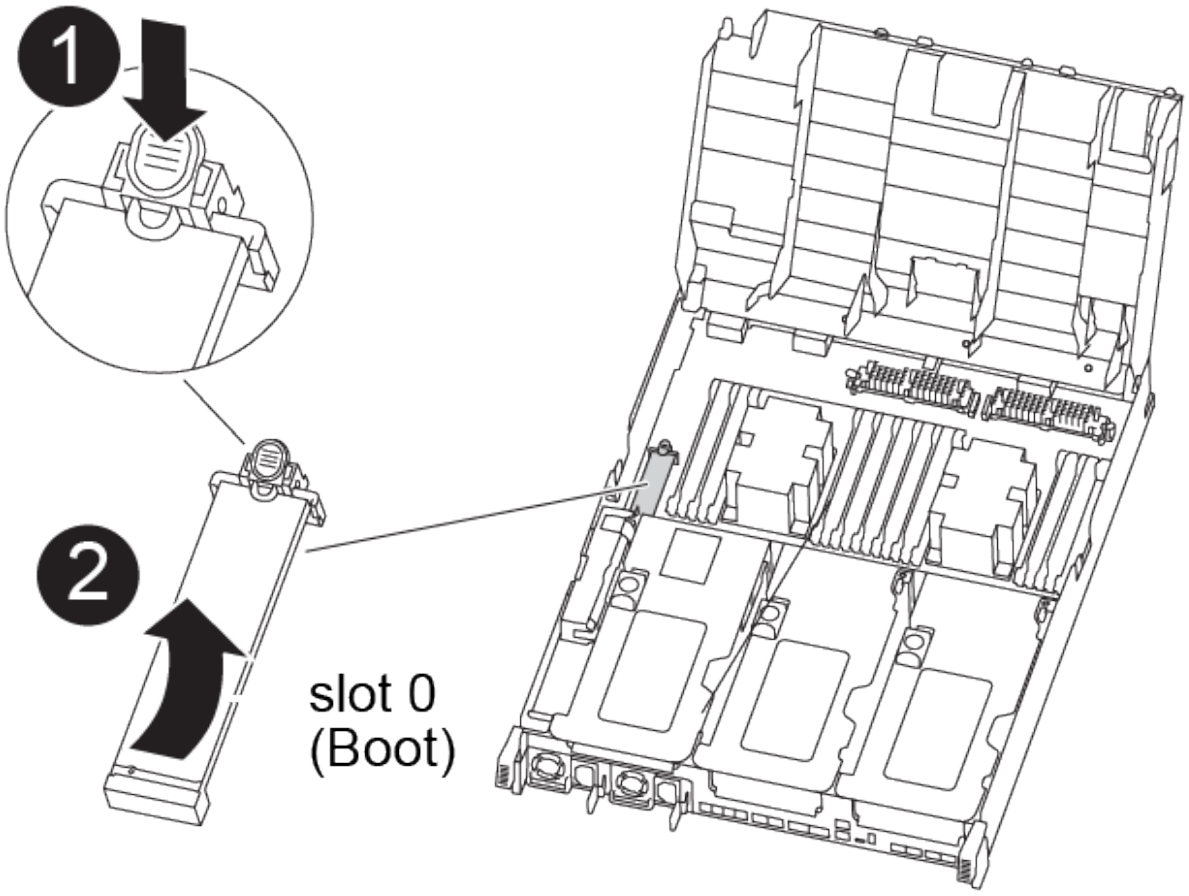

1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
b. Rotate the boot media up and gently pull the boot media out of the socket.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Lock the boot media in place:
   a. Rotate the boot media down toward the motherboard.
   b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
   c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
6. Close the air duct.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed does not have a boot image, so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the `var` file system.

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      There are two folders in the unzipped service image file:

      - `boot`
      - `efi`

   c. Copy the `efi` folder to the top directory on the USB flash drive.

      If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

      The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

   d. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.

5. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB

console port.

6. Complete the installation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
7. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then `halt` the controller to boot to LOADER.

8. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - FAS8300 and FAS8700

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `...would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore encryption - FAS8300 and FAS8700

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
********************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
********************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - FAS8300 and FAS8700

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the caching module - FAS8300 and FAS8700

You must replace the caching module in the controller module when your system registers a single AutoSupport (ASUP) message that the module has gone offline; failure to do so results in performance degradation.

> The Ver2 controller module has only one caching module socket in the FAS8300. FAS8700 does not have a VER2 controller module. The caching module functionality is not impacted by the socket removal.

- You must replace the failed component with a replacement FRU component you received from your

provider.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this tasks

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller.

Synchronize a node with the cluster

You might want to erase the contents of your caching module before replacing it.

Steps

1. Although data on the caching module is encrypted, you might want to erase any data from the impaired caching module and verify that the caching module has no data:
   a. Erase the data on the caching module: `system controller flash-cache secure-erase run -node node name localhost -device-id device_number`

      Run the `system controller flash-cache show` command if you don't know the Flash Cache device ID.

   b. Verify that the data has been erased from the caching module: `system controller flash-cache secure-erase show`
2. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`
3. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
4. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

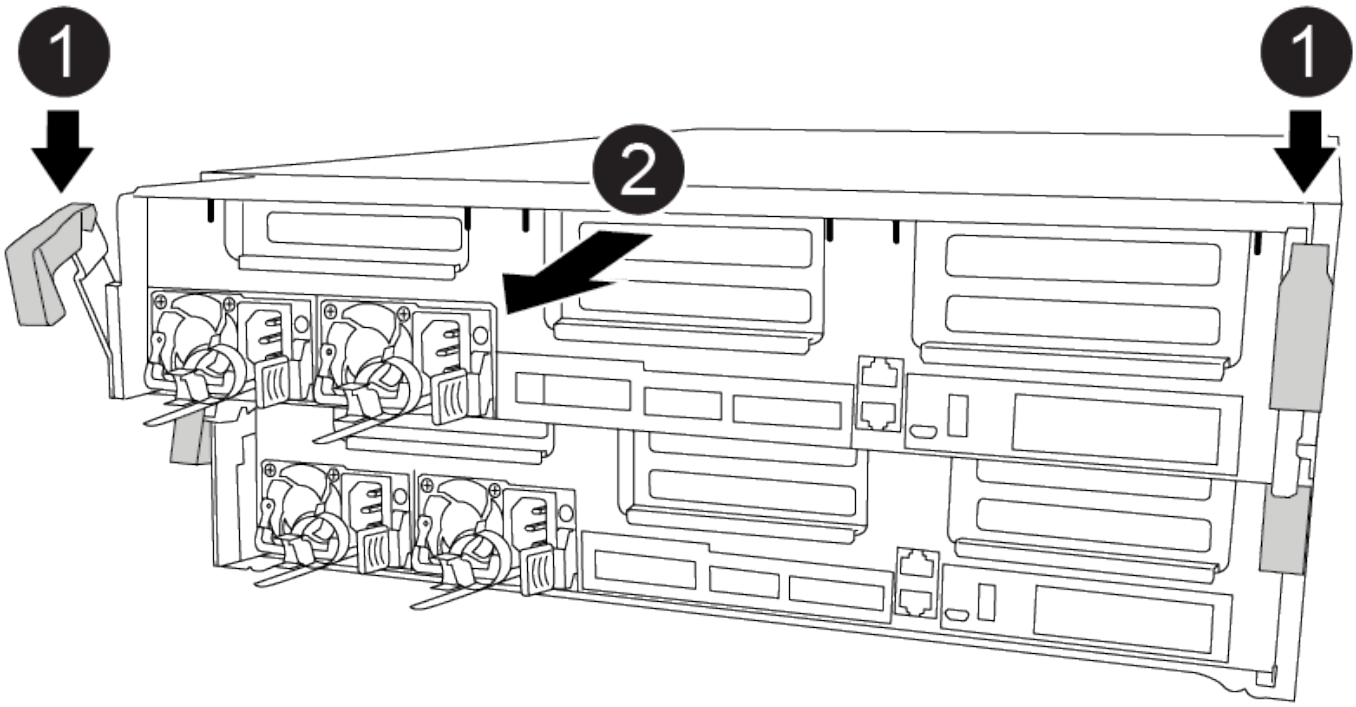

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace a caching module

To replace a caching module, referred to as the Flash Cache on the label on your controller, locate the slot inside the controller and follow the specific sequence of steps. See the FRU map on the controller module for the location of the Flash Cache.

Slot 6 is only available in FAS8300 VER2 Controller.

Your storage system must meet certain criteria depending on your situation:

- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- Although the contents of the caching module is encrypted, it is a best practice to erase the contents of the module before replacing it. For more information, see the Statement of Volatility for your system on the NetApp Support Site.

> You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

- All other components in the storage system must be functioning properly; if not, you must contact technical support.

You can use the following animation, illustration, or the written steps to replace a caching module.

Animation - Replace the caching module

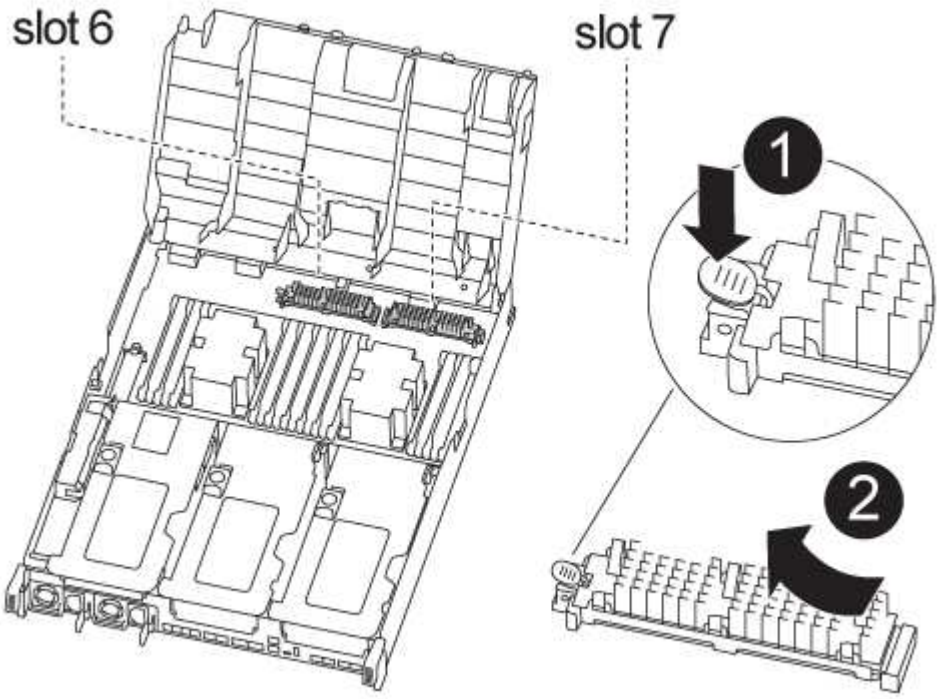

Steps

1. If you are not already grounded, properly ground yourself.
2. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
3. Using the FRU map on the controller module, locate the failed caching module and remove it:

   Depending on your configuration, there may be zero, one, or two caching modules in the controller module. Use the FRU map inside the controller module to help locate the caching module.

   a. Press the blue release tab.

The caching module end rises clear of the release tab.

b. Rotate the caching module up and slide it out of the socket.

4. Install the replacement caching module:
   a. Align the edges of the replacement caching module with the socket and gently insert it into the socket.
   b. Rotate the caching module downward toward the motherboard.
   c. Placing your finger at the end of the caching module by the blue button, firmly push down on the caching module end, and then lift the locking button to lock the caching module in place.
5. Close the air duct:
   a. Rotate the air duct down to the controller module.
   b. Slide the air duct toward the risers to lock it in place.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module

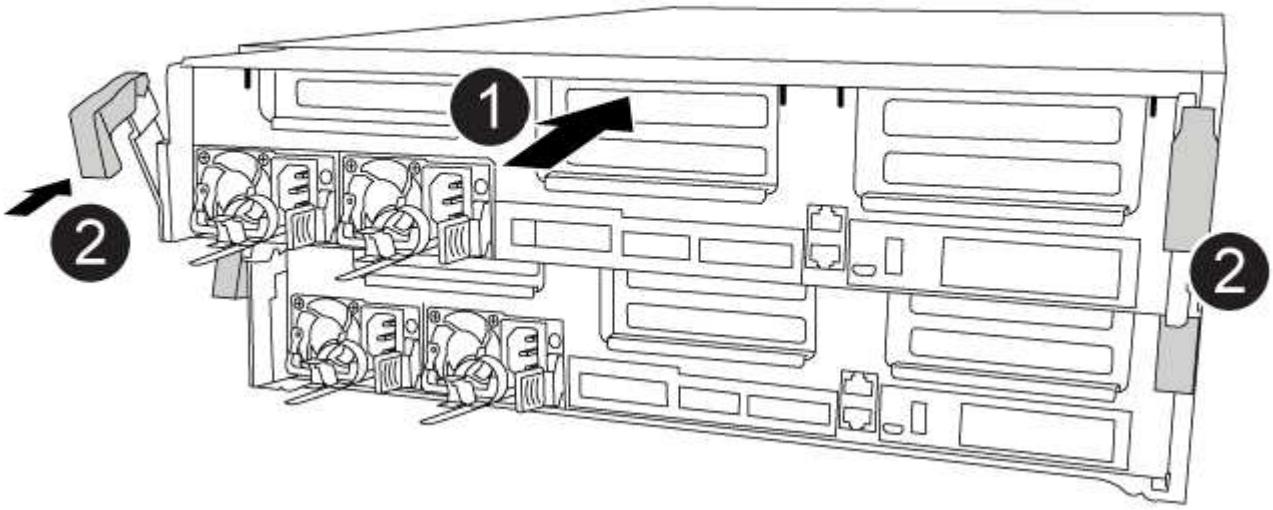

Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 7: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 8: Complete the replacement process

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - FAS8300 and FAS8700

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-controller cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - FAS8300 and FAS8700

Option 1: Most configurations

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Move and replace hardware - FAS8300 and FAS8700

Move the fans, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move the fans

To move the fan modules to the replacement chassis when replacing the chassis, you must perform a specific sequence of tasks.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and

then pulling it toward you until the bezel releases from the ball studs on the chassis frame.

3. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.

4. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

5. Set the fan module aside.
6. Repeat the preceding steps for any remaining fan modules.
7. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.
10. Repeat these steps for the remaining fan modules.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

Steps

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller

module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.
   g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - FAS8300 and FAS8700

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

a. Set the HA state for the chassis: `ha-config modify chassis *HA-state*`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`
   - `non-ha`

b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Complete the replacement process

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - FAS8300 and FAS8700

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.

- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement node* is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - FAS8300 and FAS8700

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

The Ver2 controller module has only one caching module socket (Slot 6) in the FAS8300. FAS8700 does not have a VER2 controller module. The caching module functionality is not impacted by the socket removal.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

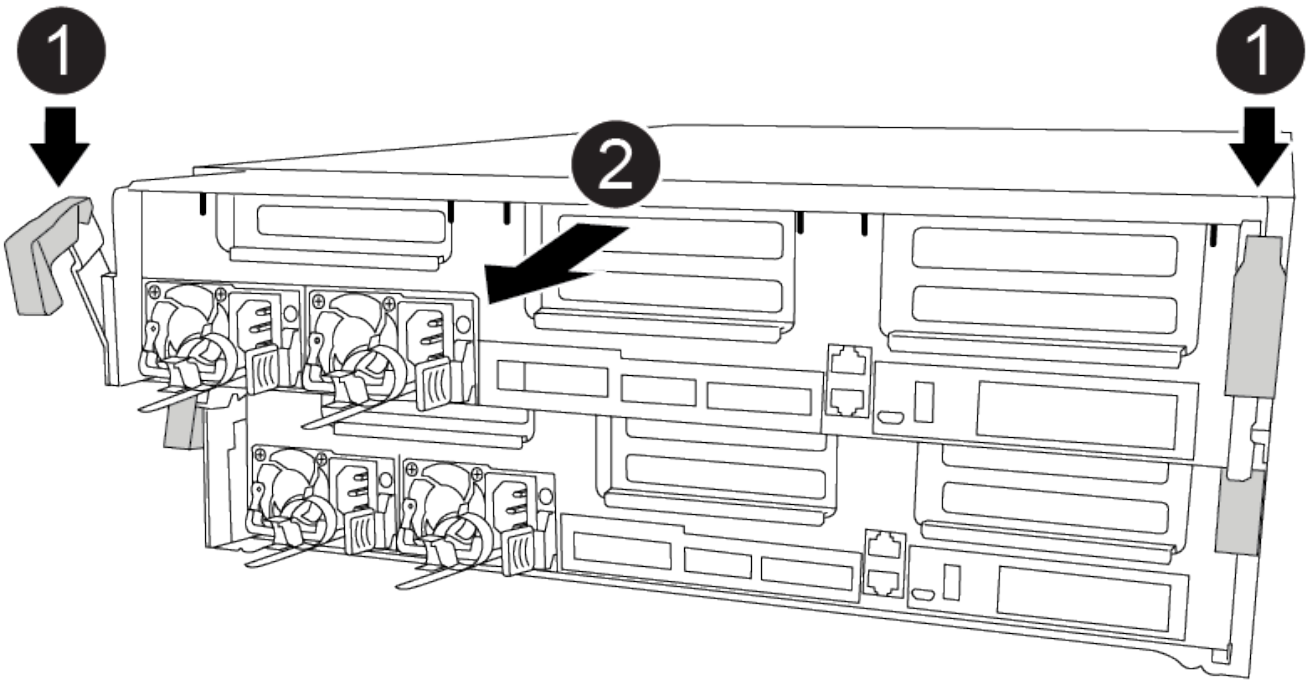

The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. On the replacement controller module, open the air duct and remove the empty risers from the controller module using the animation, illustration, or the written steps:

   Removing the empty risers from the replacement controller module

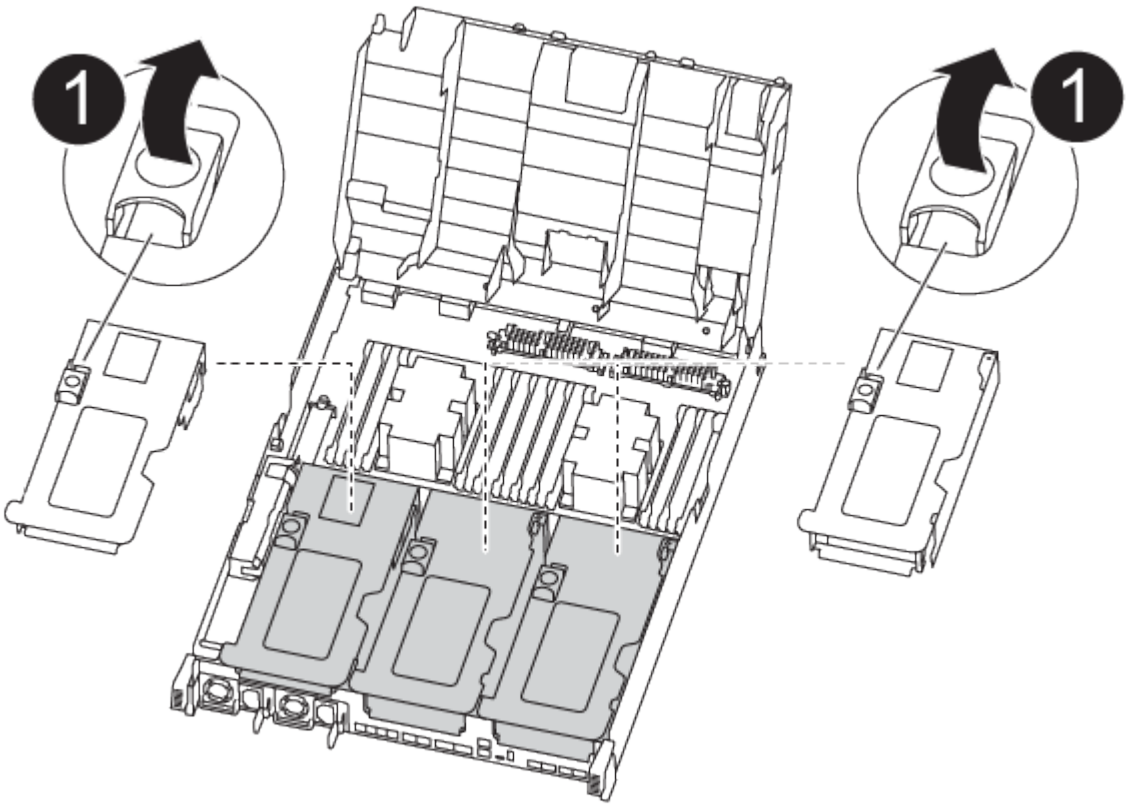

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
c. Rotate the riser locking latch on the left side of riser 1 up and toward air duct, lift the riser up, and then set it aside.
d. Repeat the previous step for the remaining risers.

Step 2: Move the power supplies

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

You can use the following animation, illustration, or the written steps to move the power supplies to the replacement controller module.

Animation - Move the power supplies

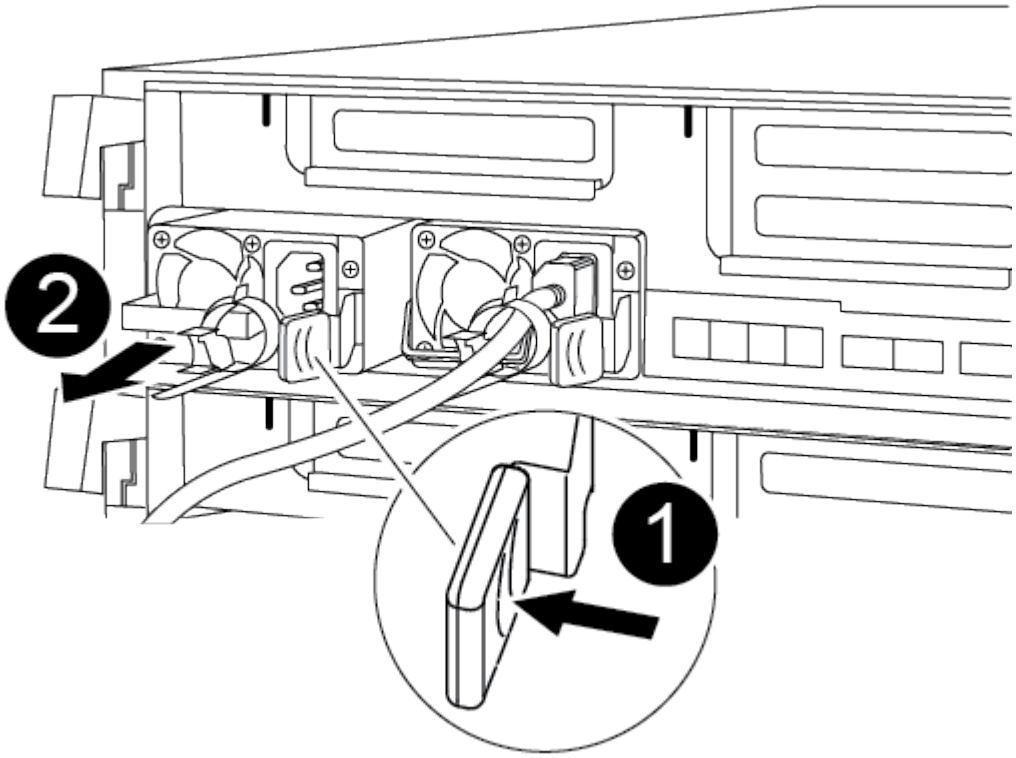

1. Remove the power supply:
   a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
   b. Press the blue locking tab to release the power supply from the chassis.
   c. Using both hands, pull the power supply out of the chassis, and then set it aside.
2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

4. Repeat the preceding steps for any remaining power supplies.

Step 3: Move the NVDIMM battery

To move the NVDIMM battery from the impaired controller module to the replacement controller module, you must perform a specific sequence of steps.

You can use the following animation, illustration, or the written steps to move the NVDIMM battery from the impaired controller module to the replacement controller module.

Animation - Move the NVDIMM battery

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.

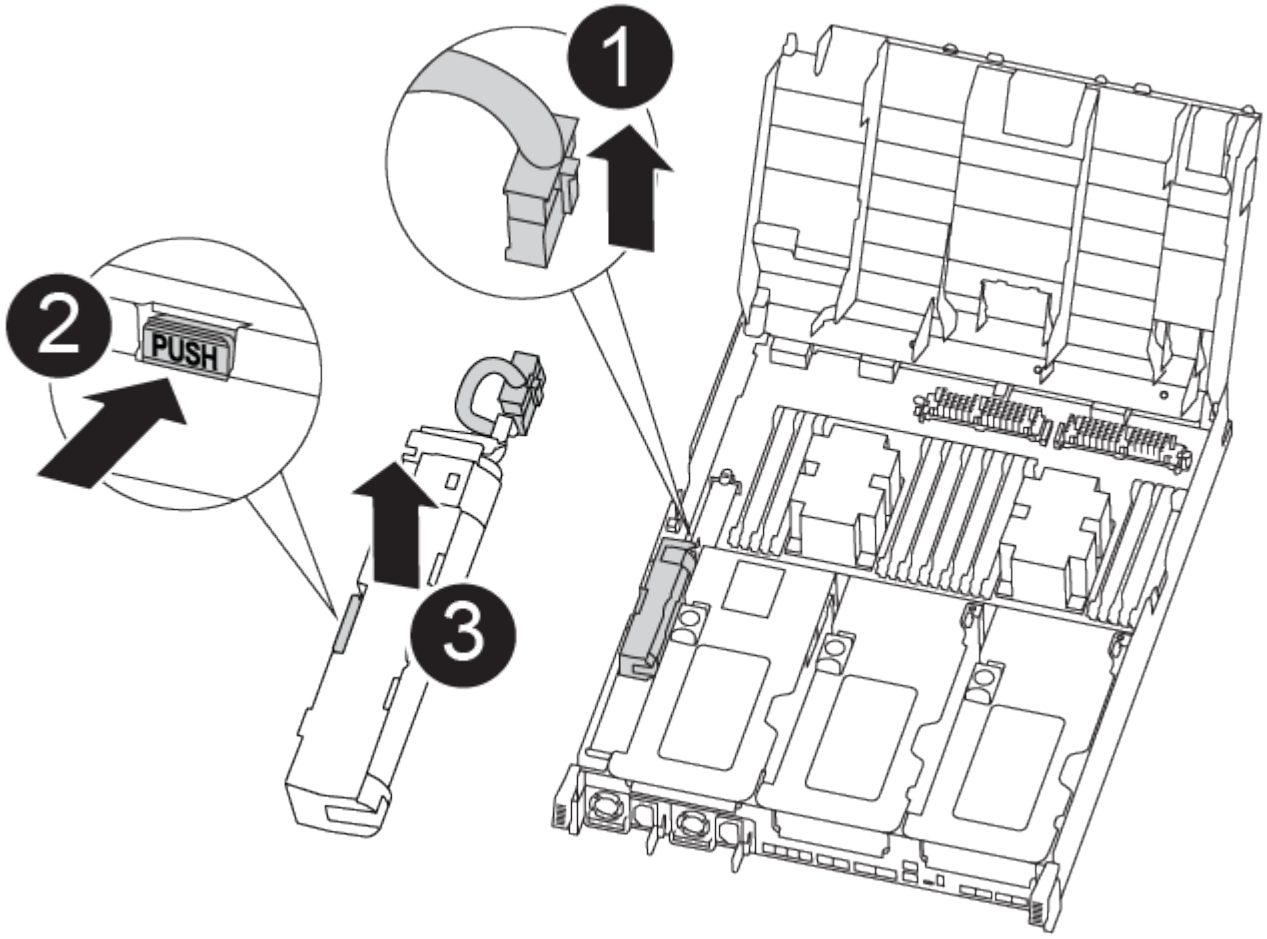

1. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
2. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
3. Move the battery to the replacement controller module.
4. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.

> Do not plug the battery cable back into the motherboard until instructed to do so.

Step 4: Move the boot media

You must locate the boot media, and then follow the directions to remove it from the impaired controller module and insert it into the replacement controller module.

You can use the following animation, illustration, or the written steps to move the boot media from the impaired controller module to the replacement controller module.

Animation - Move the boot media

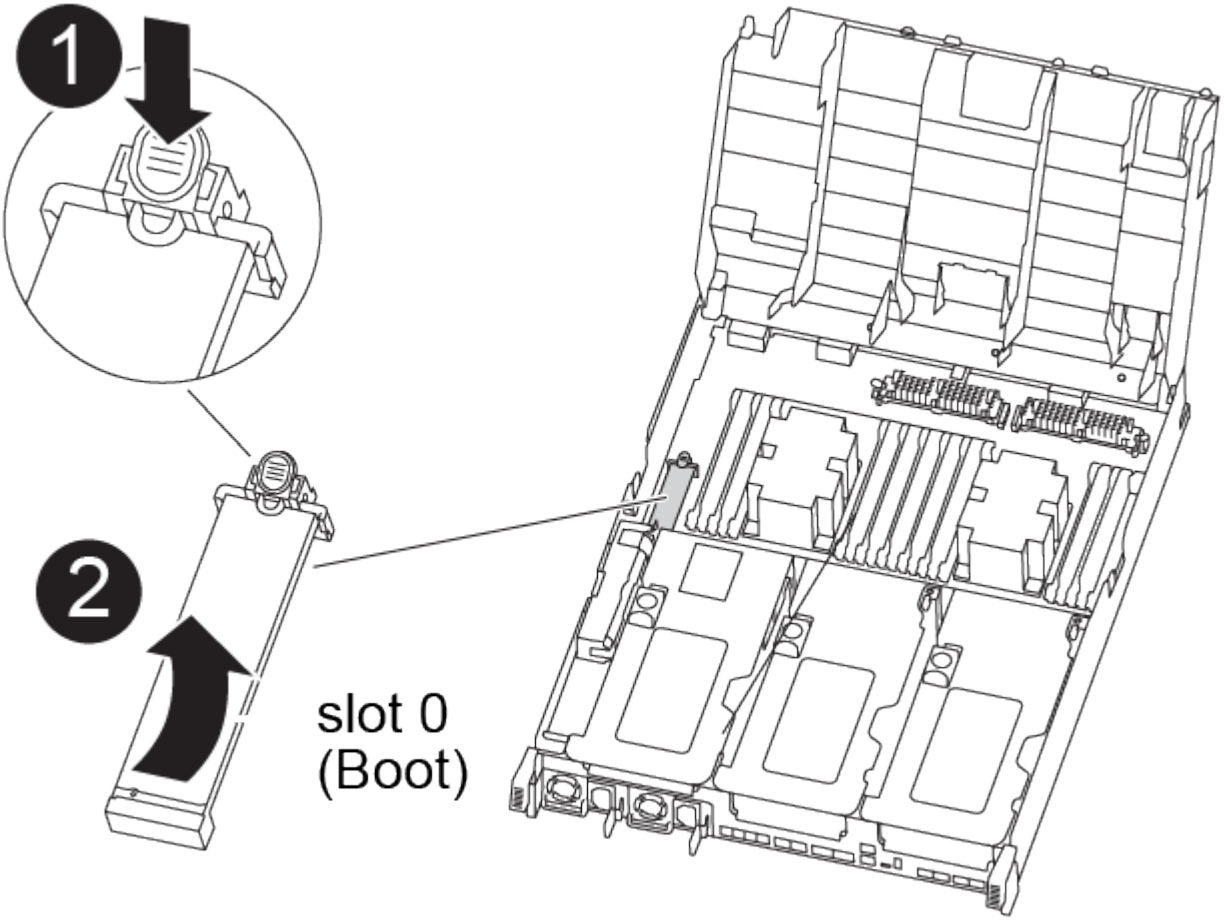

1. Locate and remove the boot media from the controller module:
   a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
   b. Rotate the boot media up and gently pull the boot media out of the socket.
2. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
3. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

4. Lock the boot media in place:
   a. Rotate the boot media down toward the motherboard.
   b. Press the blue locking button so that it is in the open position.
   c. Placing your fingers at the end of the boot media by the blue button, firmly push down on the boot media end to engage the blue locking button.

Step 5: Move the PCIe risers and mezzanine card

As part of the controller replacement process, you must move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

You can use the following animations, illustrations, the FUR map on the system, or the written steps to move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

> You do not have to remove the PCIe cards from the risers. Transfer the risers, with the PCIe cards still installed, to the replacement controller module.

Moving PCIe riser 1 and 2 (left and middle risers):

Animation - Move PCI risers 1 and 2

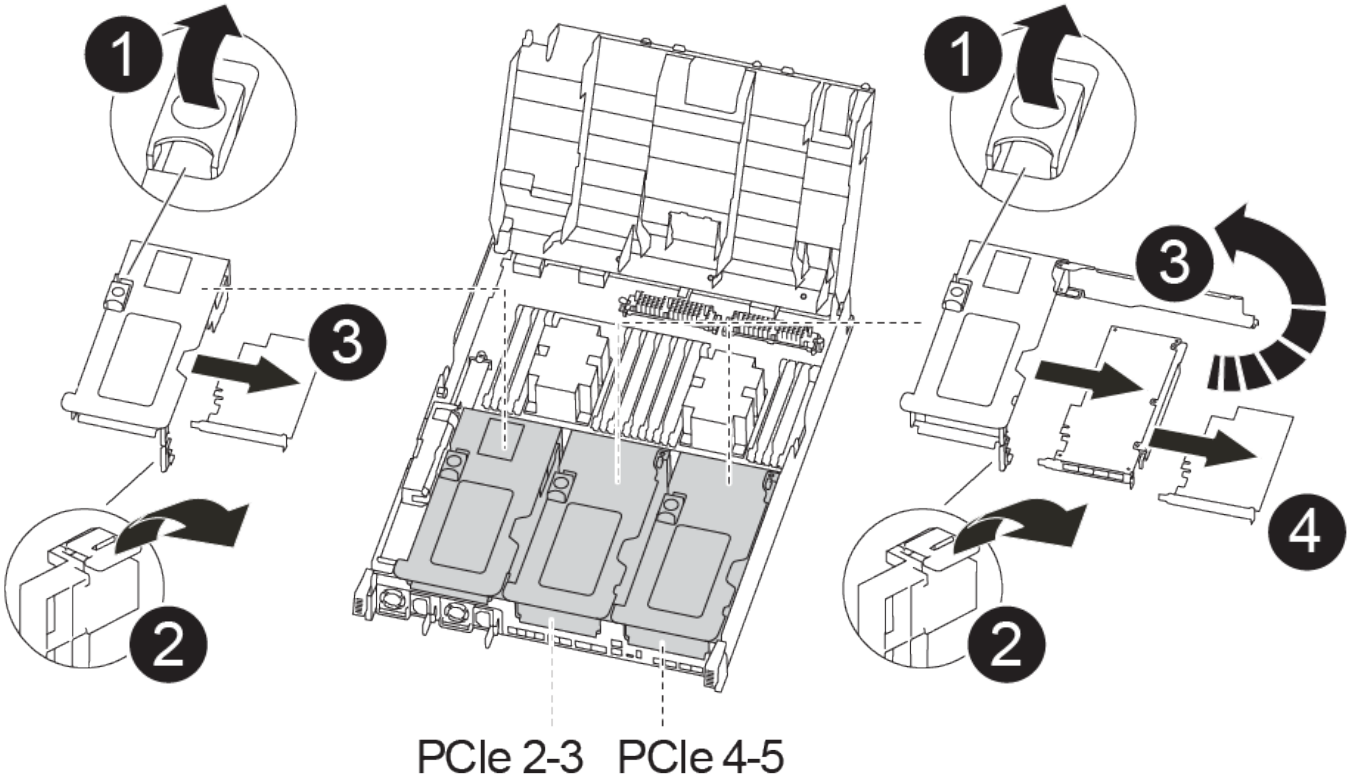

Moving the mezzanine card and riser 3 (right riser):

Animation - Move the mezzanine card and riser 3

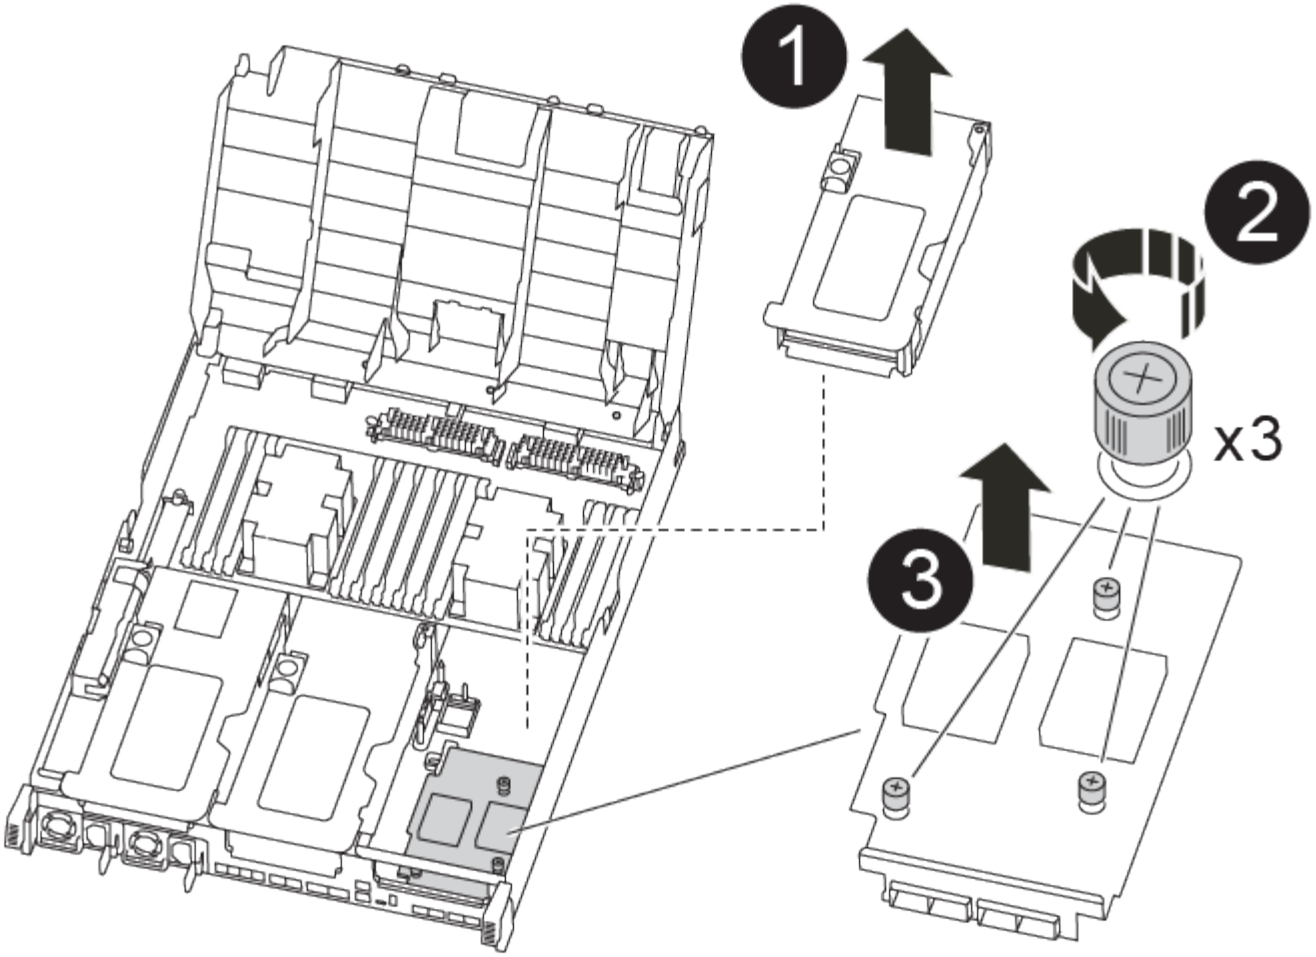

1. Move PCIe risers one and two from the impaired controller module to the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then move it to the replacement controller module.
   d. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins, push the riser squarely into the socket on the motherboard, and then rotate the latch down flush with the sheet metal on the riser.
   e. Repeat this step for riser number 2.
2. Remove riser number 3, remove the mezzanine card, and install both into the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then set it aside on a stable, flat surface.
   d. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket, and then move it to the replacement controller module.

e. Install the mezzanine in the replacement controller and secure it with the thumbscrews.

f. Install the third riser in the replacement controller module.

Step 6: Move caching modules

You must move the caching modules from the impaired controller modules to the replacement controller module when replacing a controller module.

> The Ver2 controller module has only one caching module socket in the FAS8300. FAS8700 does not have a VER2 controller module. The caching module functionality is not impacted by the socket removal.

You can use the following animation, illustration, or the written steps to move caching modules to the new controller module.

Animation - Move the caching modules

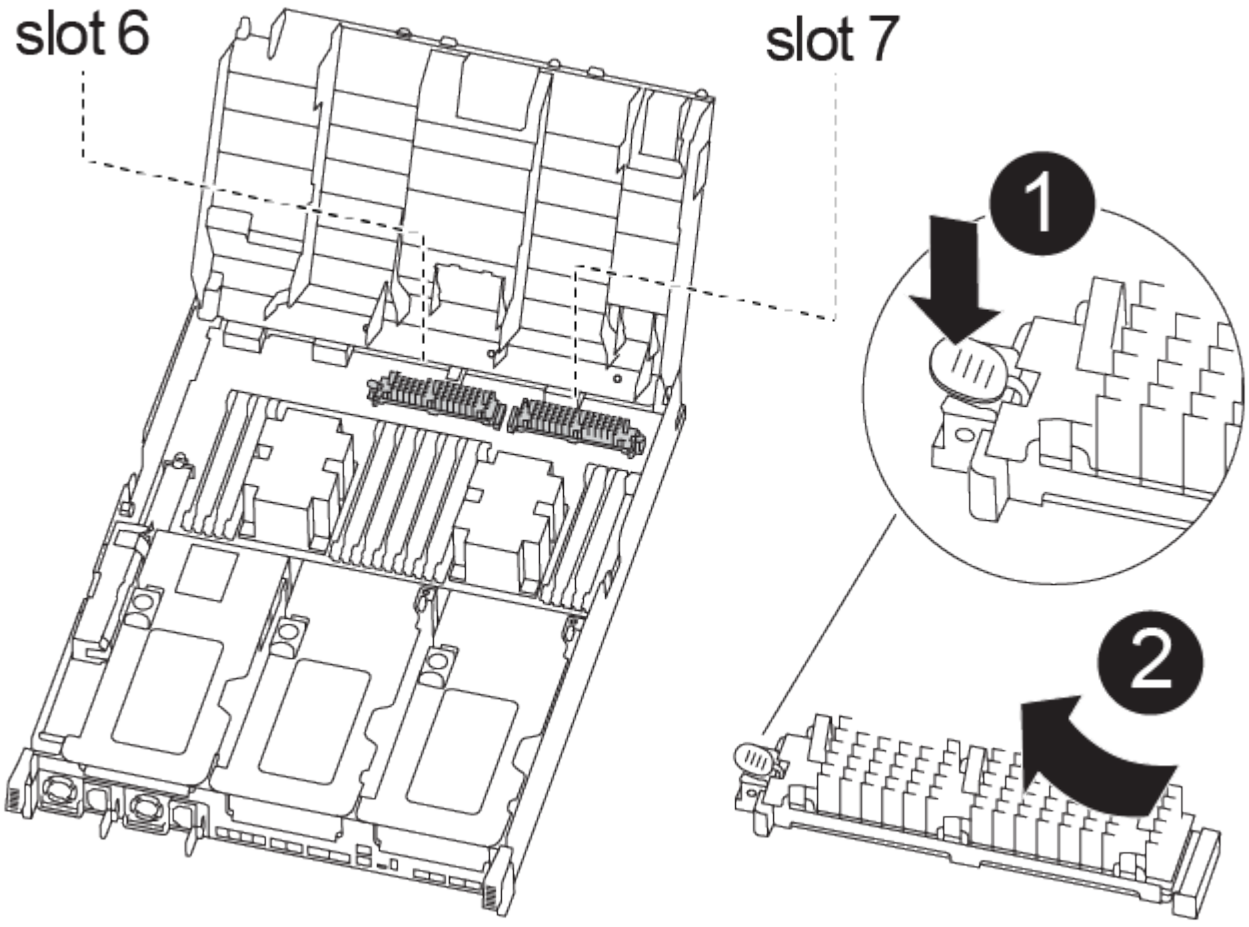

1. If you are not already grounded, properly ground yourself.
2. Move the caching modules from the impaired controller module to the replacement controller module:
   a. Press the blue release tab at the end of the caching module, rotate the module up, and then remove the module from the socket.
   b. Move the caching module to the same socket on the replacement controller module.

c. Align the edges of the caching module with the socket and gently insert the module as far into the socket as it will go.

d. Rotate the caching module downward toward the motherboard.

e. Placing your finger at the end of the caching module by the blue button, firmly push down on the caching module end, and then lift the locking button to lock the caching module in place.

Step 7: Move the DIMMs

You need to locate the DIMMs, and then move them from the impaired controller module to the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

You can use the following animation, illustration, or the written steps to move the DIMMs from the impaired controller module to the replacement controller module.

Animation - Move the DIMMs

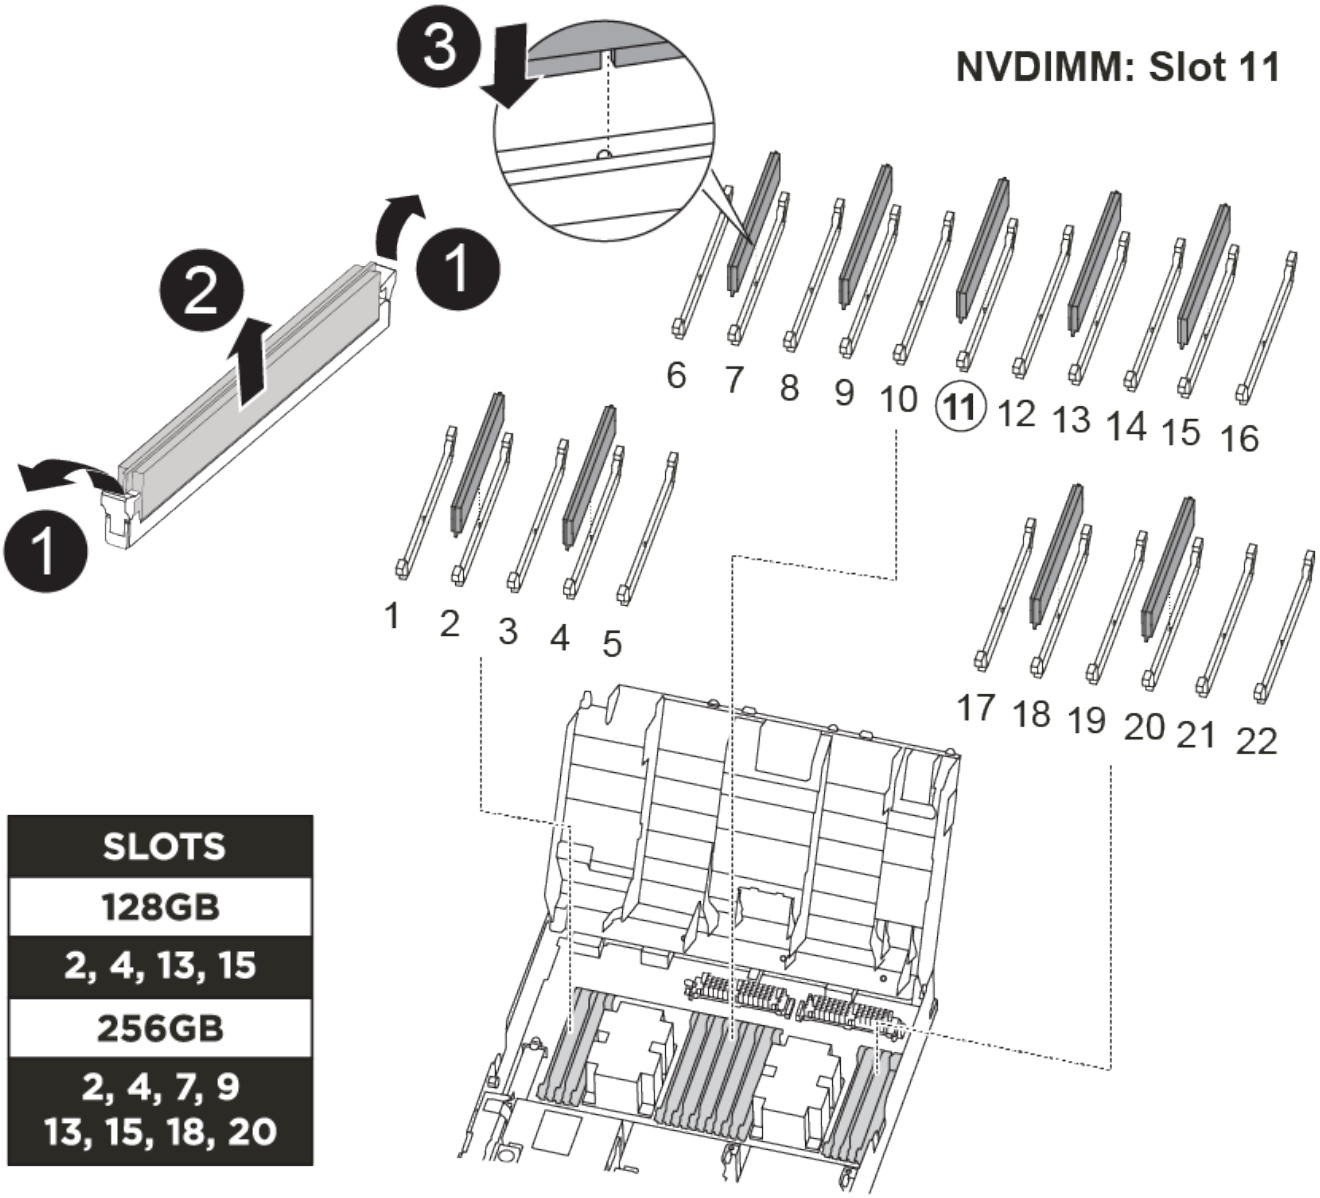

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Verify that the NVDIMM battery is not plugged into the new controller module.
4. Move the DIMMs from the impaired controller module to the replacement controller module:

> Make sure that you install the each DIMM into the same slot it occupied in the impaired controller module.

a. Eject the DIMM from its slot by slowly pushing apart the DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

> Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

b. Locate the corresponding DIMM slot on the replacement controller module.

c. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the DIMM squarely into the socket.

   The DIMMs fit tightly in the socket, but should go in easily. If not, realign the DIMM with the socket and reinsert it.

d. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

e. Repeat these substeps for the remaining DIMMs.

5. Plug the NVDIMM battery into the motherboard.

   Make sure that the plug locks down onto the controller module.

Step 8: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

You can use the following animation, illustration, or the written steps to install the replacement controller module in the chassis.

Animation - Install the controller module

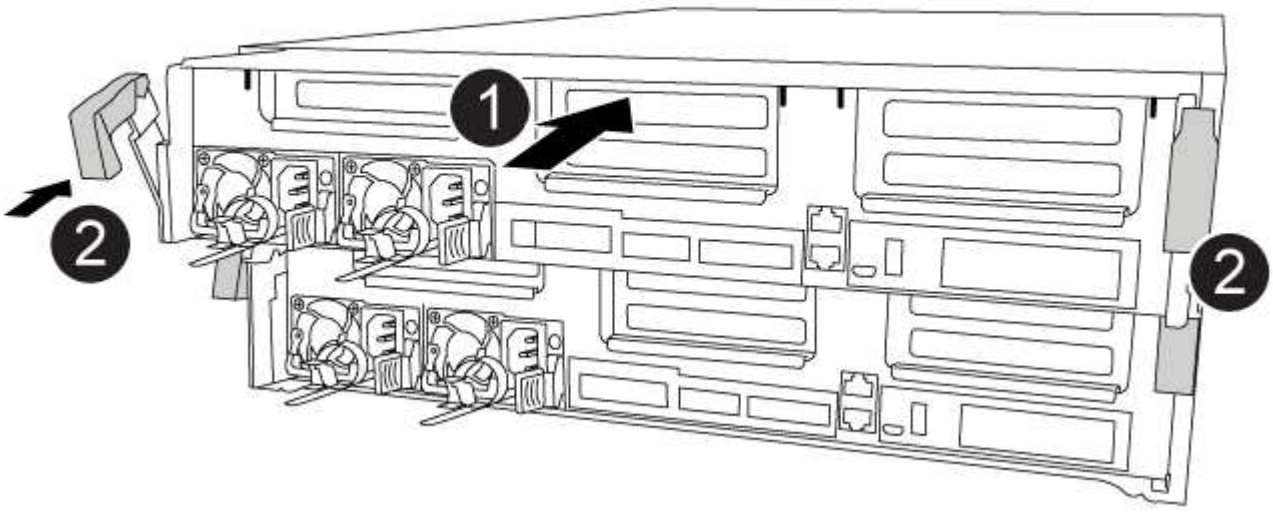

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.
   g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

Restore and verify the system configuration - FAS8300 and FAS8700

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.

2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for `HA-state` can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - FAS8300 and FAS8700

You must complete a series of tasks before restoring your system to full operation.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.

a. Download and install Config Advisor.

b. Enter the information for the target system, and then click Collect Data.

c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.

d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks.

You must use the correct procedure for your configuration:

Controller redundancy	Then use this procedure…
HA pair	Option 1: Verify the system ID change on an HA system]
Two-node MetroCluster configuration	Option 2: Manually reassign the system ID on systems in a two-node MetroCluster configuration

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node                Partner         Possible     State Description
------------      ------------      --------
-------------------------------------
node1               node2           false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2               node1           -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration Guide for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed

on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Option 2: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `metrocluster node show -fields node-systemid`,dr-partner-systemid`

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

```
dr-group-id cluster         node                 node-systemid dr-
partner-systemid
 ----------- -------------------- ------------------- ------------
------------------
 1           Cluster_A            Node_A_1             536872914
118073209
 1           Cluster_B            Node_B_1             118073209
536872914
 2 entries were displayed.
```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

```
Local System ID: 118065481
    ...
    ...
```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                POOL   SERIAL NUMBER  HOME
-------    -------------        -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:

a. Change to the advanced privilege level: `set -privilege advanced`

   You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

b. Verify that the coredumps are saved: `system node run -node local-node-name partner savecore`

   If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node local-node-name partner savecore -s command`.</info>.

c. Return to the admin privilege level: `set -privilege admin`

7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
8. Boot the *replacement* node: `boot_ontap`
9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`
10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`

c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - FAS8300 and FAS8700

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`

b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS8300 and FAS8700

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

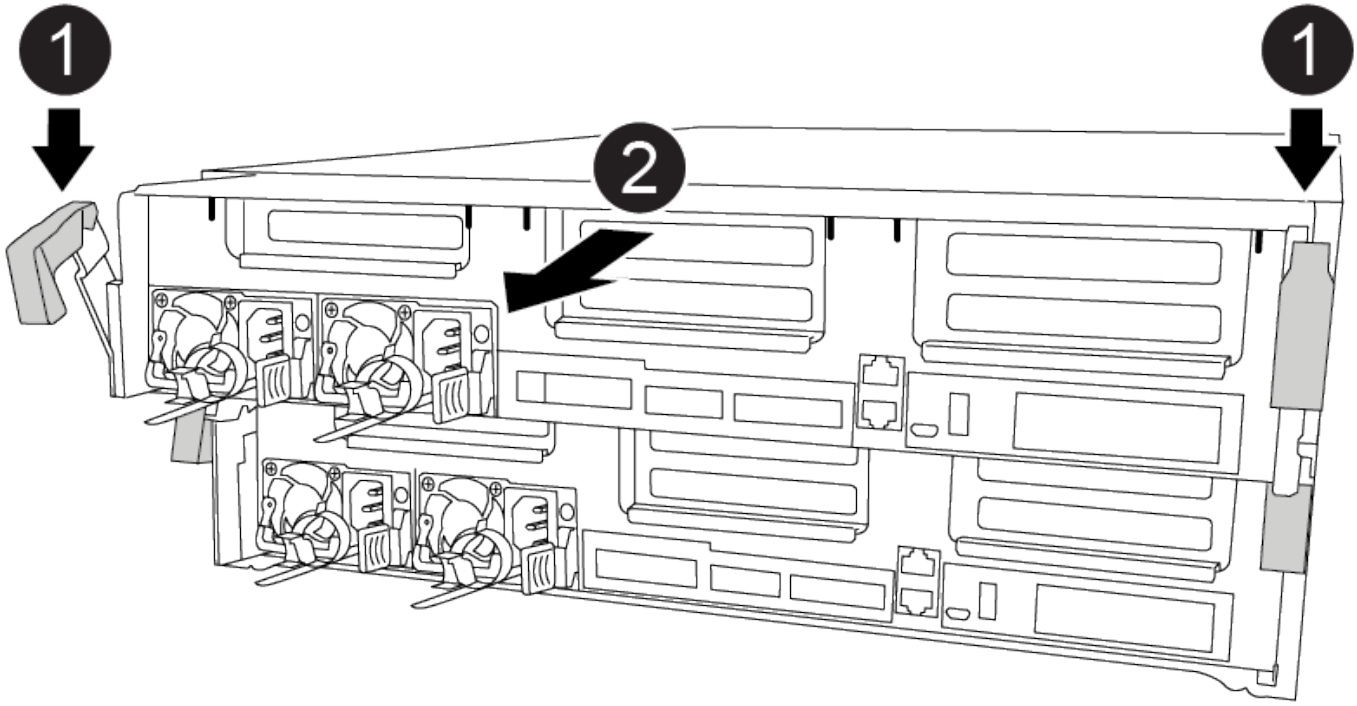

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Place the controller module on a stable, flat surface.

Step 3: Replace system DIMMs

Replacing a system DIMM involves identifying the target DIMM through the associated error message, locating the target DIMM using the FRU map on the air duct, and then replacing the DIMM.

You can use the following animation, illustration, or the written steps to replace a system DIMM.

> The animation and illustration shows empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace a system DIMM

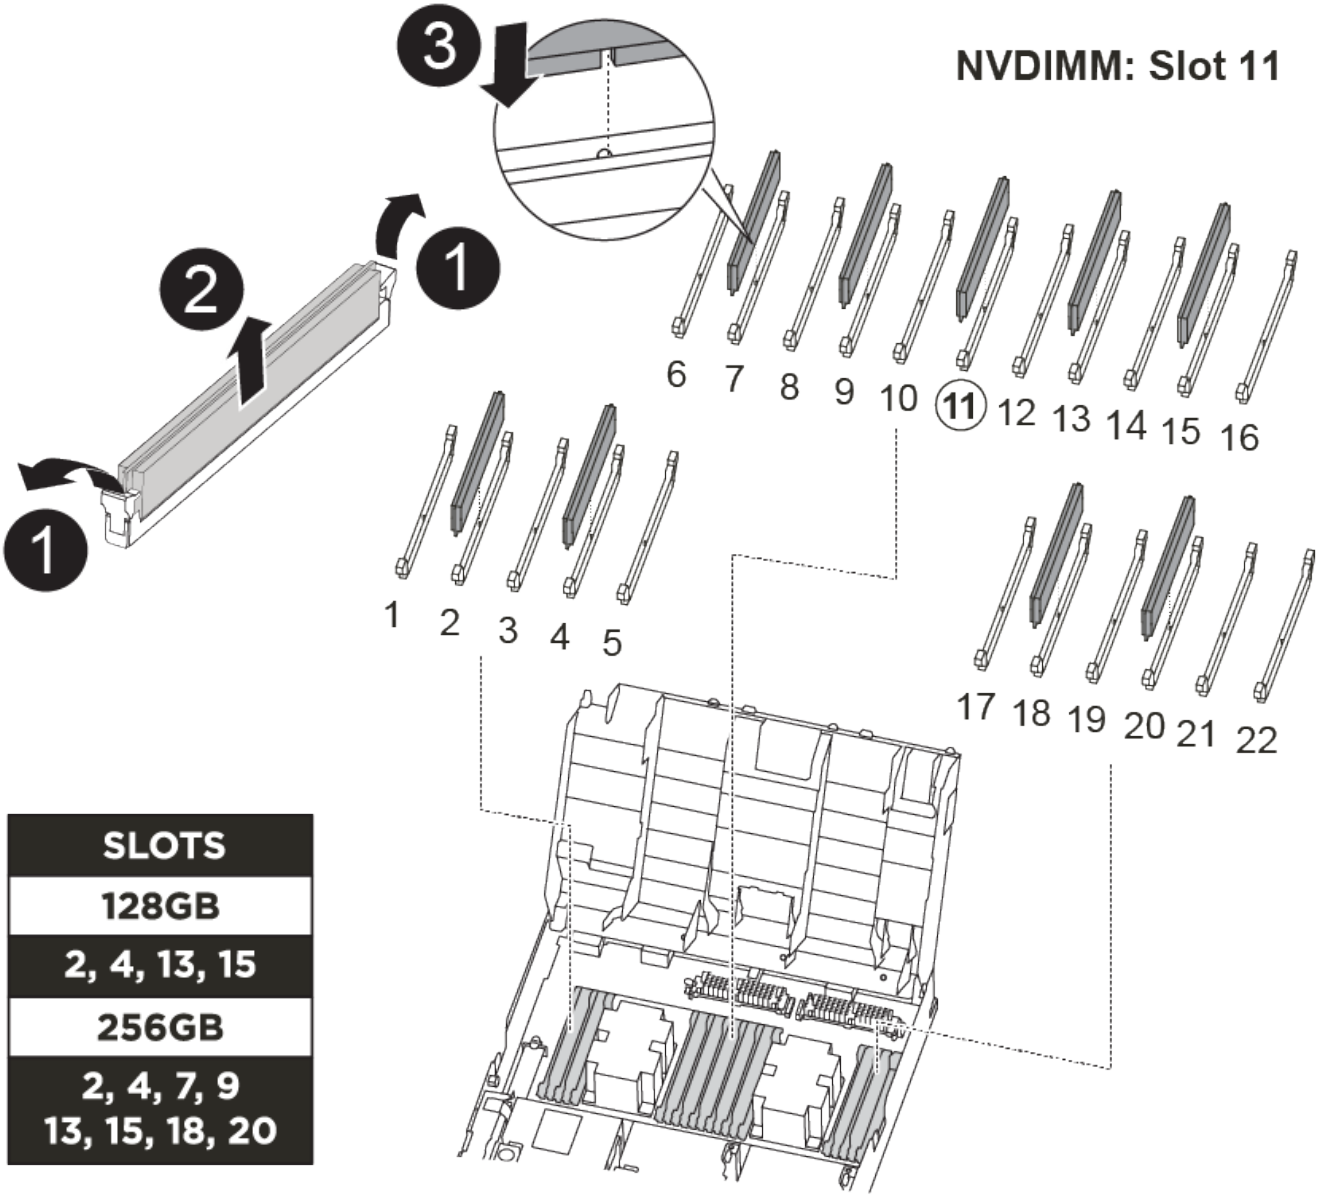

The number and location of DIMMS in your system depends on the model of your system. Refer to FRU map on the air duct for more information.

- If you have a FAS8300 system, the system DIMMs are located in sockets 2, 4, 13, and 15.
- If you have a FAS8700 system, the system DIMMs are located in slots 2, 4, 7, 9, 13, 15, 18, and 20.

- The NVDIMM is located in slot 11.

Steps

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
4. Eject the DIMM from its socket by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the socket.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis.

You can use the following animation, drawing, or the written steps to install the controller module in the chassis.

Animation - Install the controller module

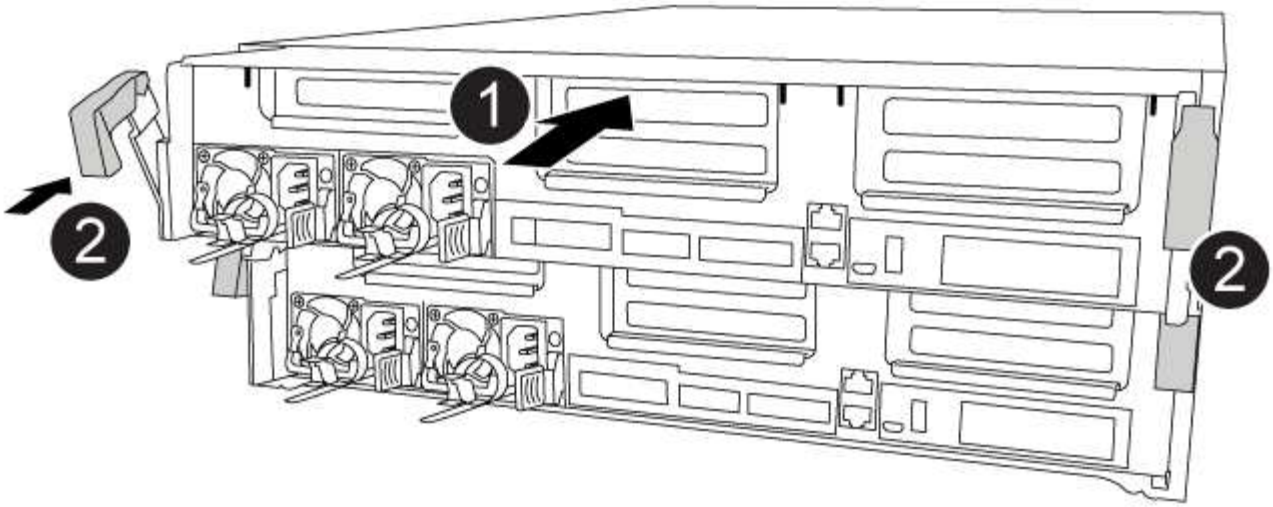

Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a fan module - FAS8300 and FAS8700

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

You can use the following animation, illustration, or the written steps to hot-swap a fan module.

Animation - Replace a fan

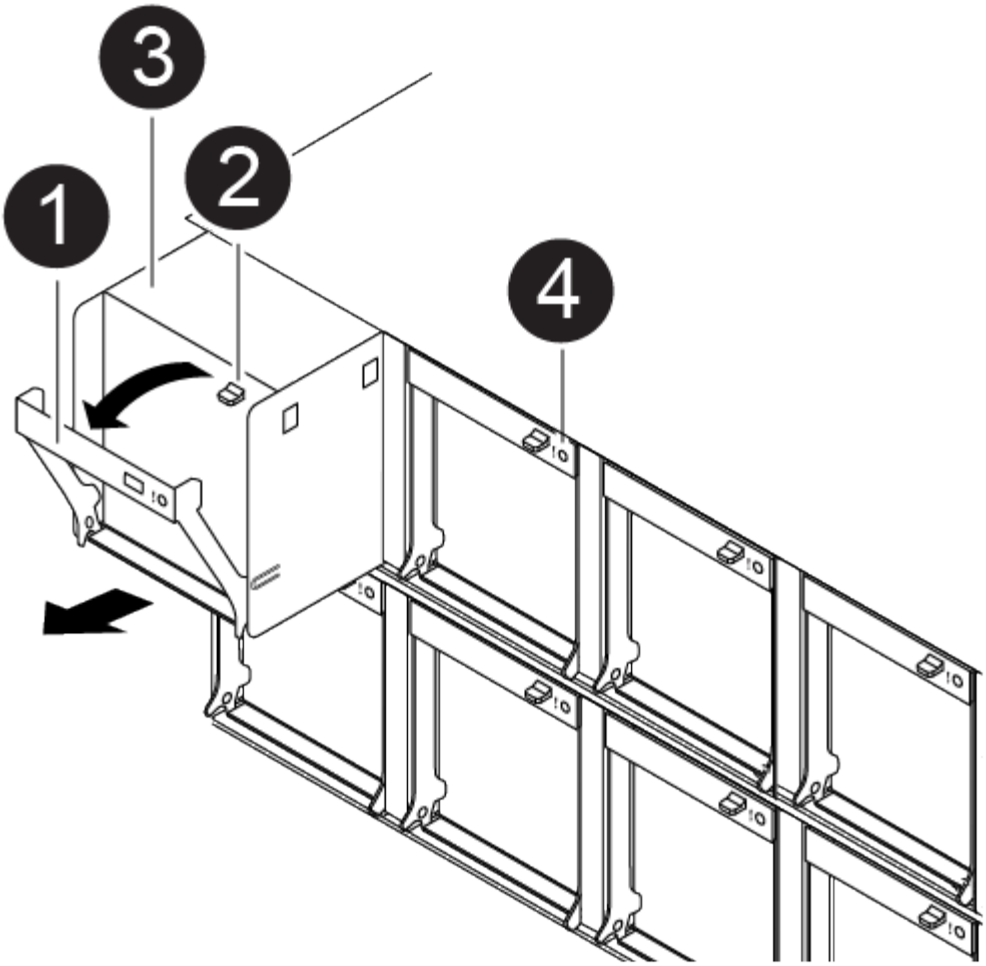

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.

5. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

6. Set the fan module aside.
7. Insert the replacement fan module into the chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The Attention LED should not be lit after the fan is seated and has spun up to operational speed.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an NVDIMM - FAS8300 and FAS8700

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following , illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

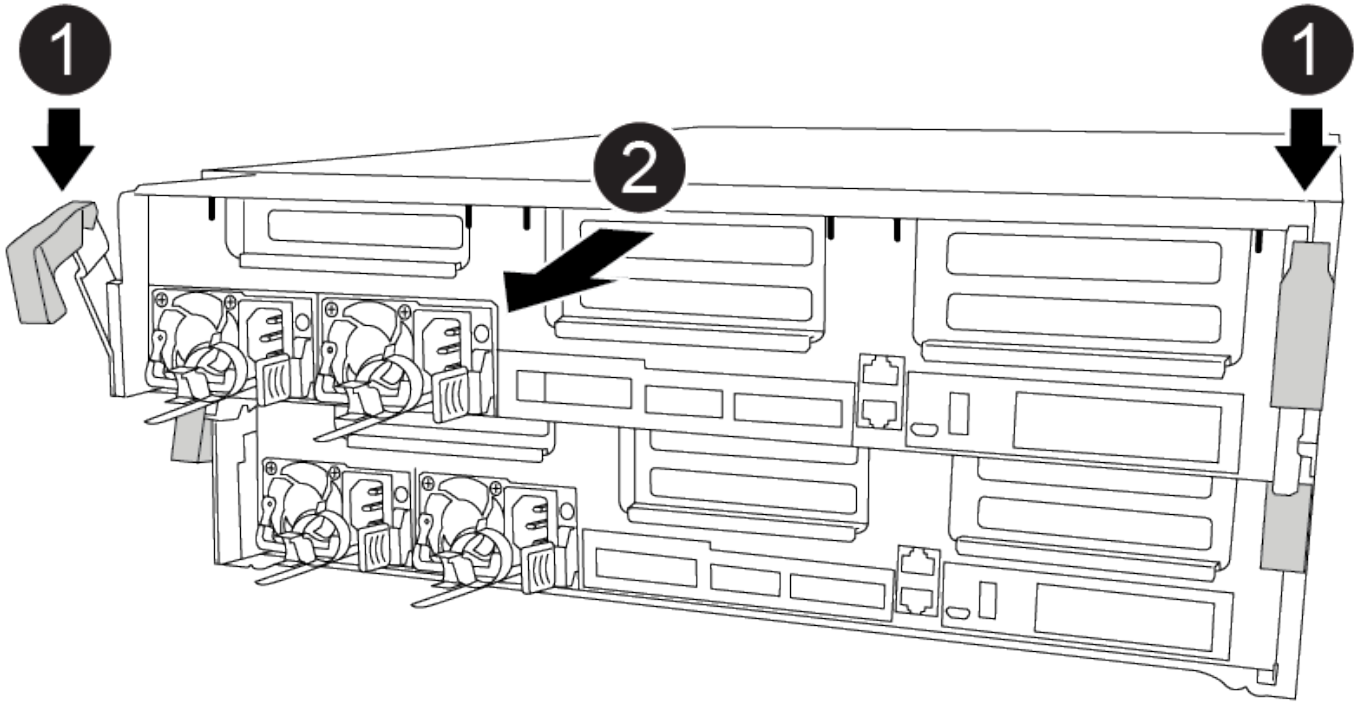

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the FRU map on top of the air duct the FRU Map on the top of the slot 1 riser.

- The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.
- Although the contents of the NVDIMM is encrypted, it is a best practice to erase the contents of the NVDIMM before replacing it. For more information, see the Statement of Volatility on the NetApp Support Site.

> You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the NVDIMM.

> The animation and illustration show empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace the NVDIMM

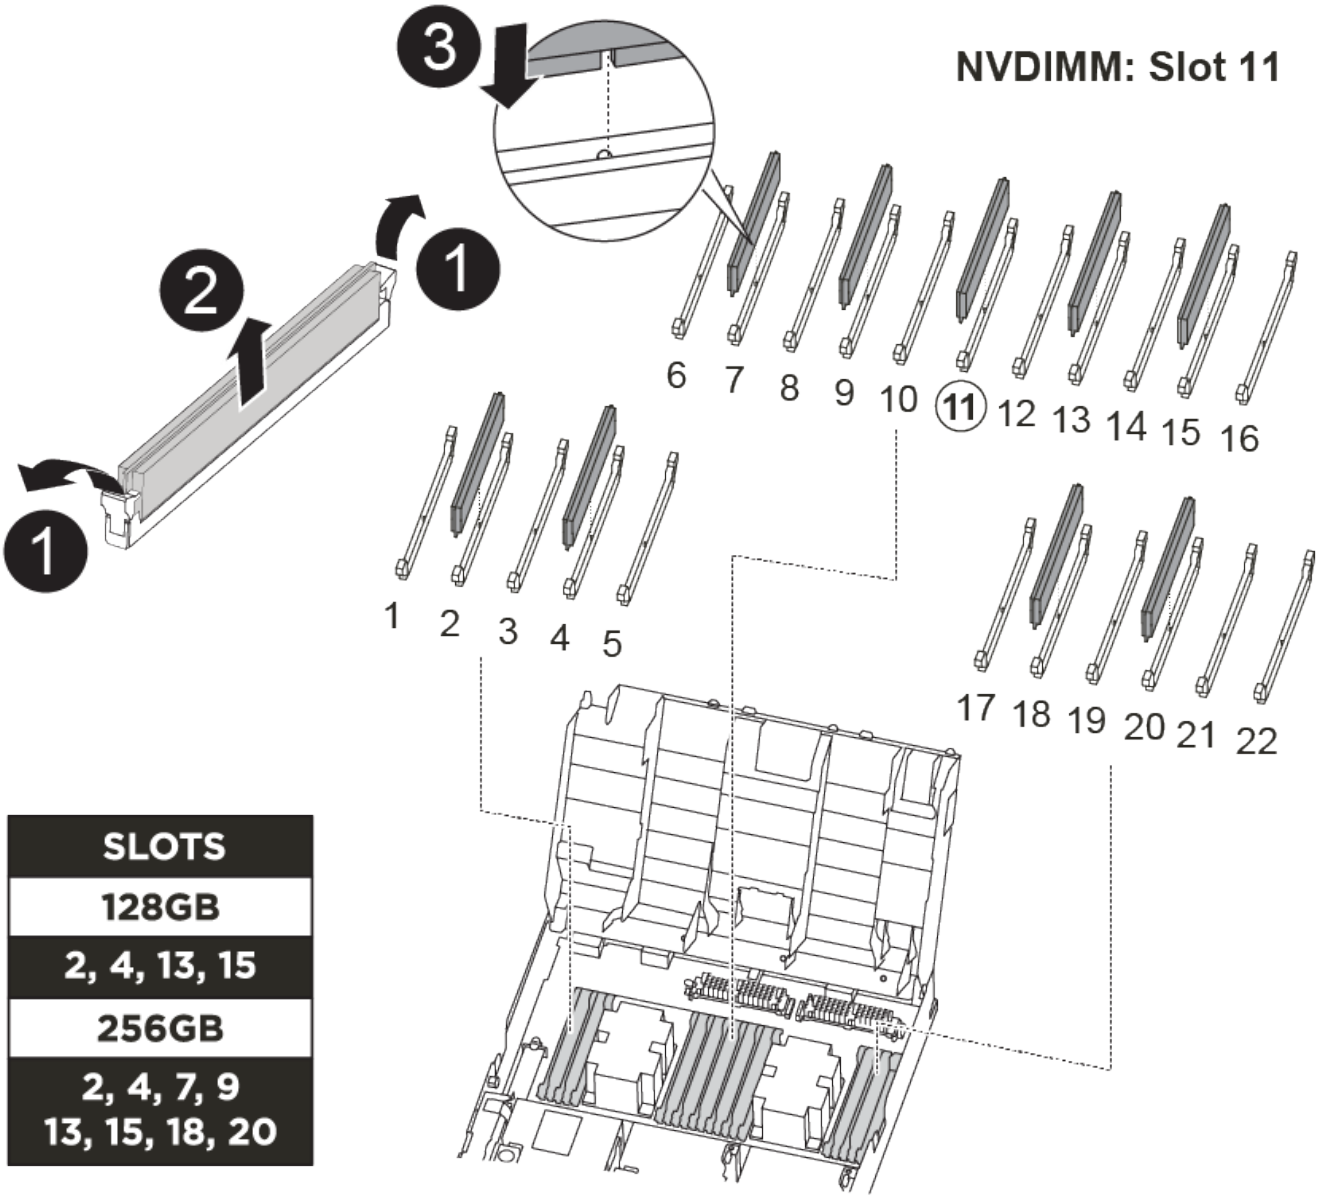

Steps

1. Open the air duct and then locate the NVDIMM in slot 11 on your controller module.

   The NVDIMM looks significantly different than system DIMMs.

2. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

3. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

4. Locate the slot where you are installing the NVDIMM.
5. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
7. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module

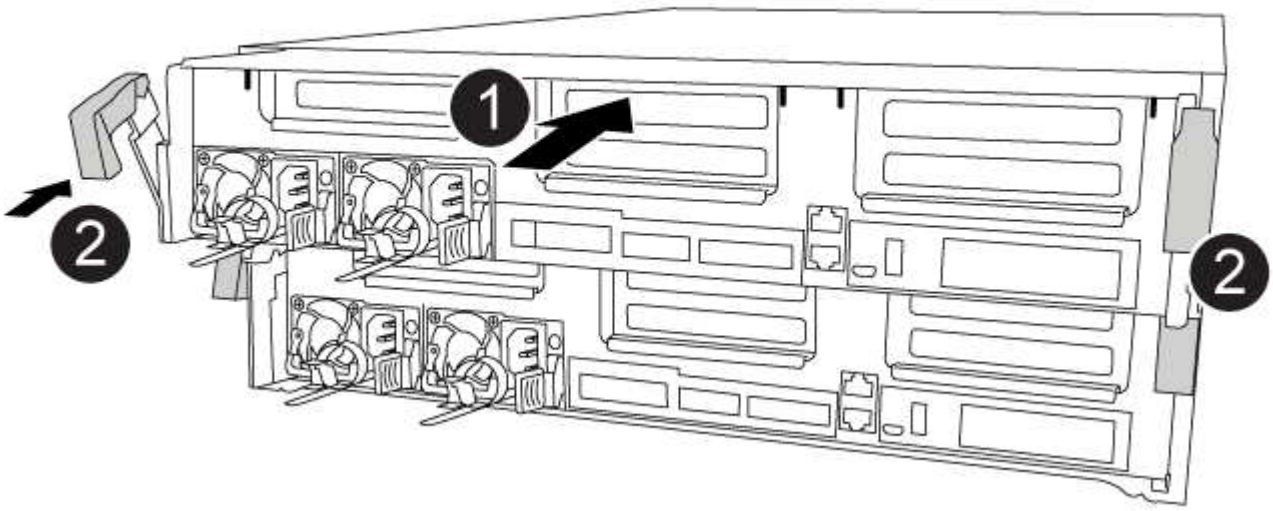

Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` `*impaired_node_name*`
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - FAS8300 and FAS8700

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

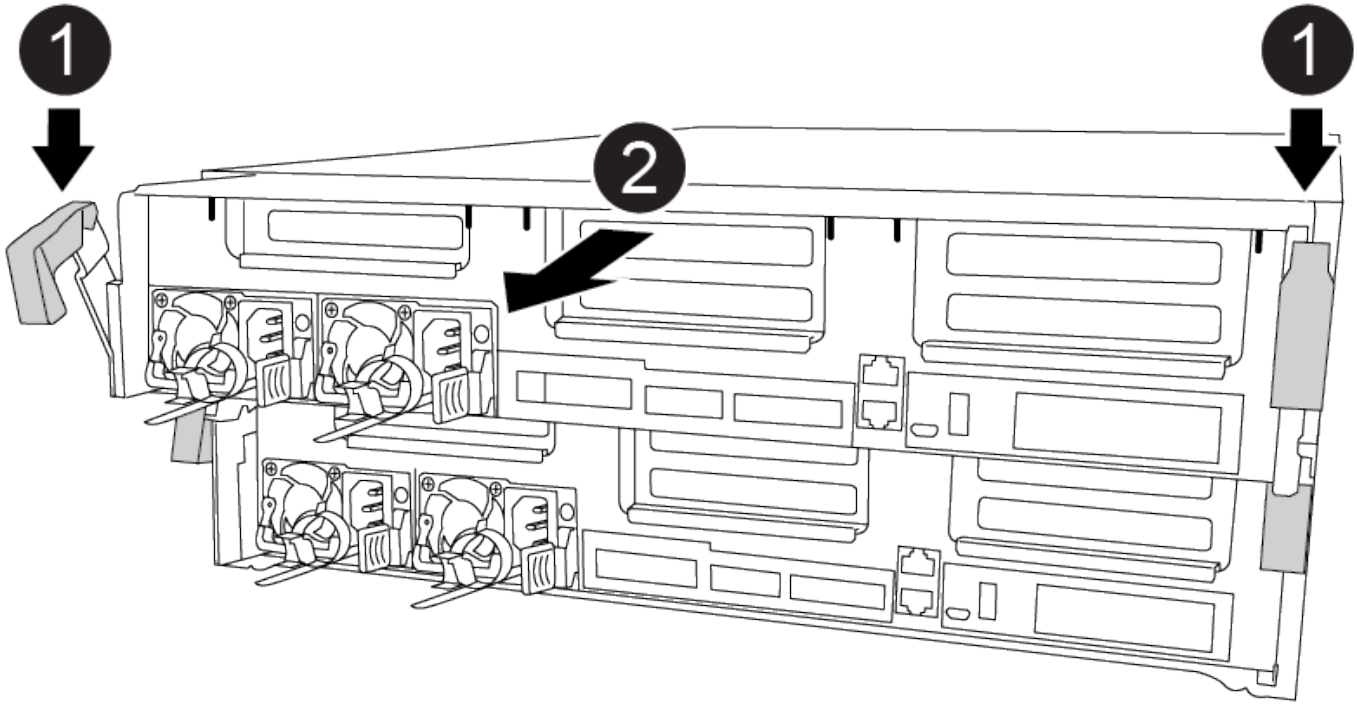

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module. See the FRU map inside the controller module to locate the NVDIMM battery.

The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.

You can use the following animation, illustration, or the written steps to replace the NVDIMM battery.

Animation - Replace the NVDIMM battery

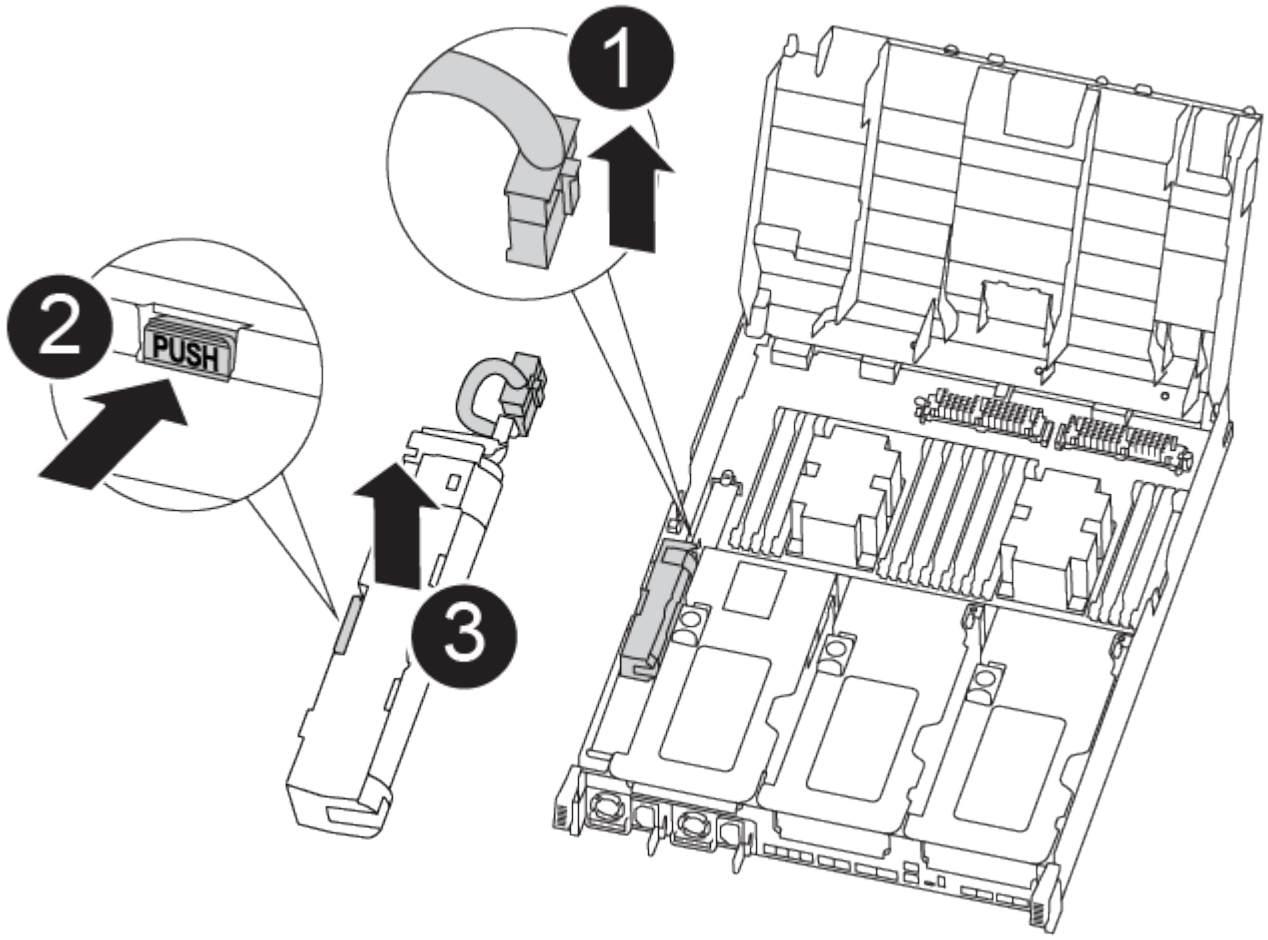

Steps

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.
3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder

and controller module.

5. Remove the replacement battery from its package.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.
7. Plug the battery plug back into the controller module, and then close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module

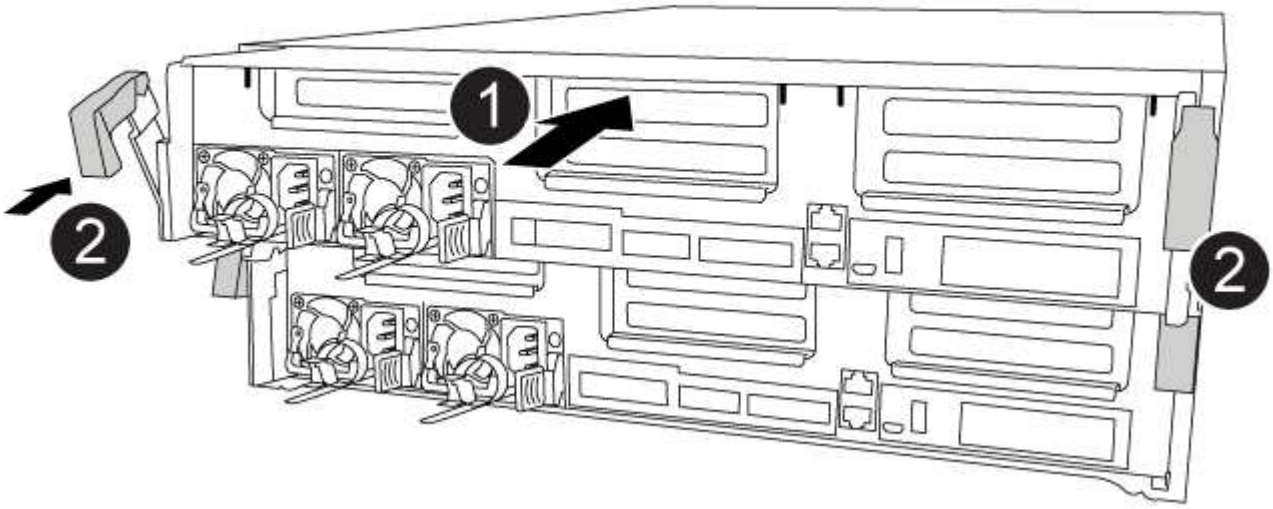

Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe or mezzanine card - FAS8300 and FAS8700

To replace a PCIe or mezzanine card, you must disconnect the cables and any SFP and QSFP modules from the cards, replace the failed PCIe or mezzanine card, and then recable the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

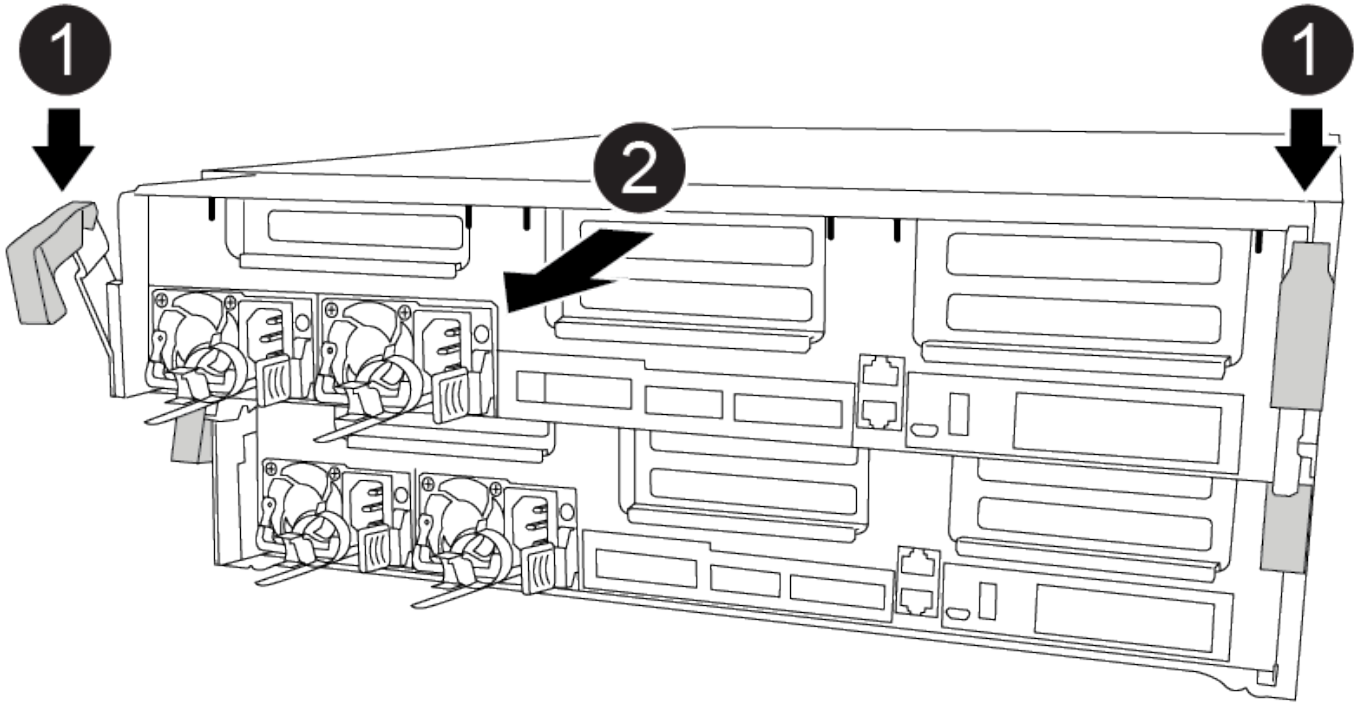

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Place the controller module on a stable, flat surface.

Step 3: Replace a PCIe card

To replace a PCIe card, you must locate the failed PCIe card, remove the riser that contains the card from the controller module, replace the card, and then reinstall the PCIe riser in the controller module.

You can use the following animation, illustration, or the written steps to replace a PCIe card.

Animation - Replace a PCIe card

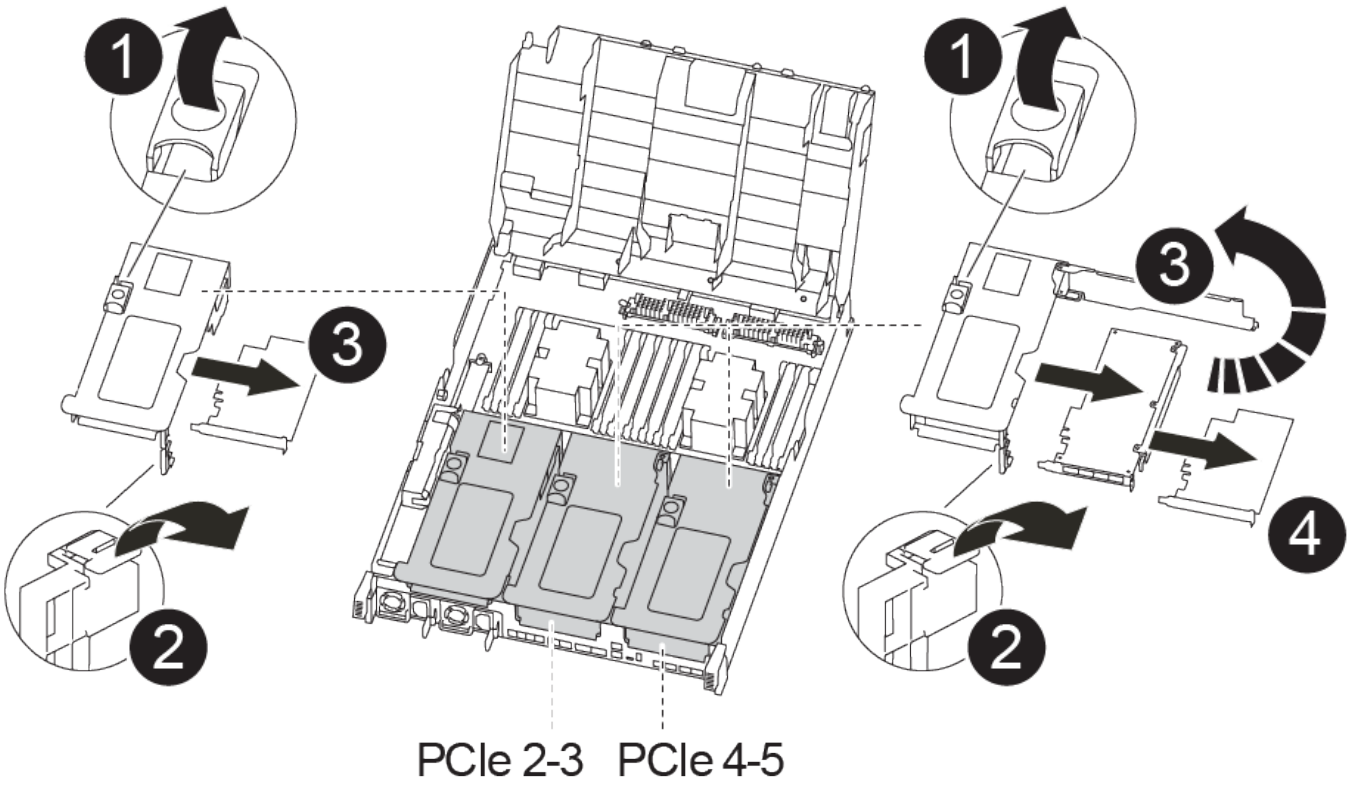

Steps

1. Remove the riser containing the card to be replaced:
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up straight up and set it aside on a stable flat surface,
2. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. For risers 2 and 3 only, swing the side panel up.
   d. Remove the PCIe card from the riser by gently pushing up on the bracket and lift the card straight out of the socket.
3. Install the replacement PCIe card in the riser by aligning the card with the socket, press the card into the

socket and then close the side panel on the riser, if present.

Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The PCIe card must be fully and evenly seated in the slot.

> If you are installing a card in the bottom slot and cannot see the card socket well, remove the top card so that you can see the card socket, install the card, and then reinstall the card you removed from the top slot.

4. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 4: Replace the mezzanine card

The mezzanine card is located under riser number 3 (slots 4 and 5). You must remove that riser to access the mezzanine card, replace the mezzanine card, and then reinstall riser number 3. See the FRU map on the controller module for more information.

You can use the following animation, illustration, or the written steps to replace the mezzanine card.

Animation - Replace the mezzanine card

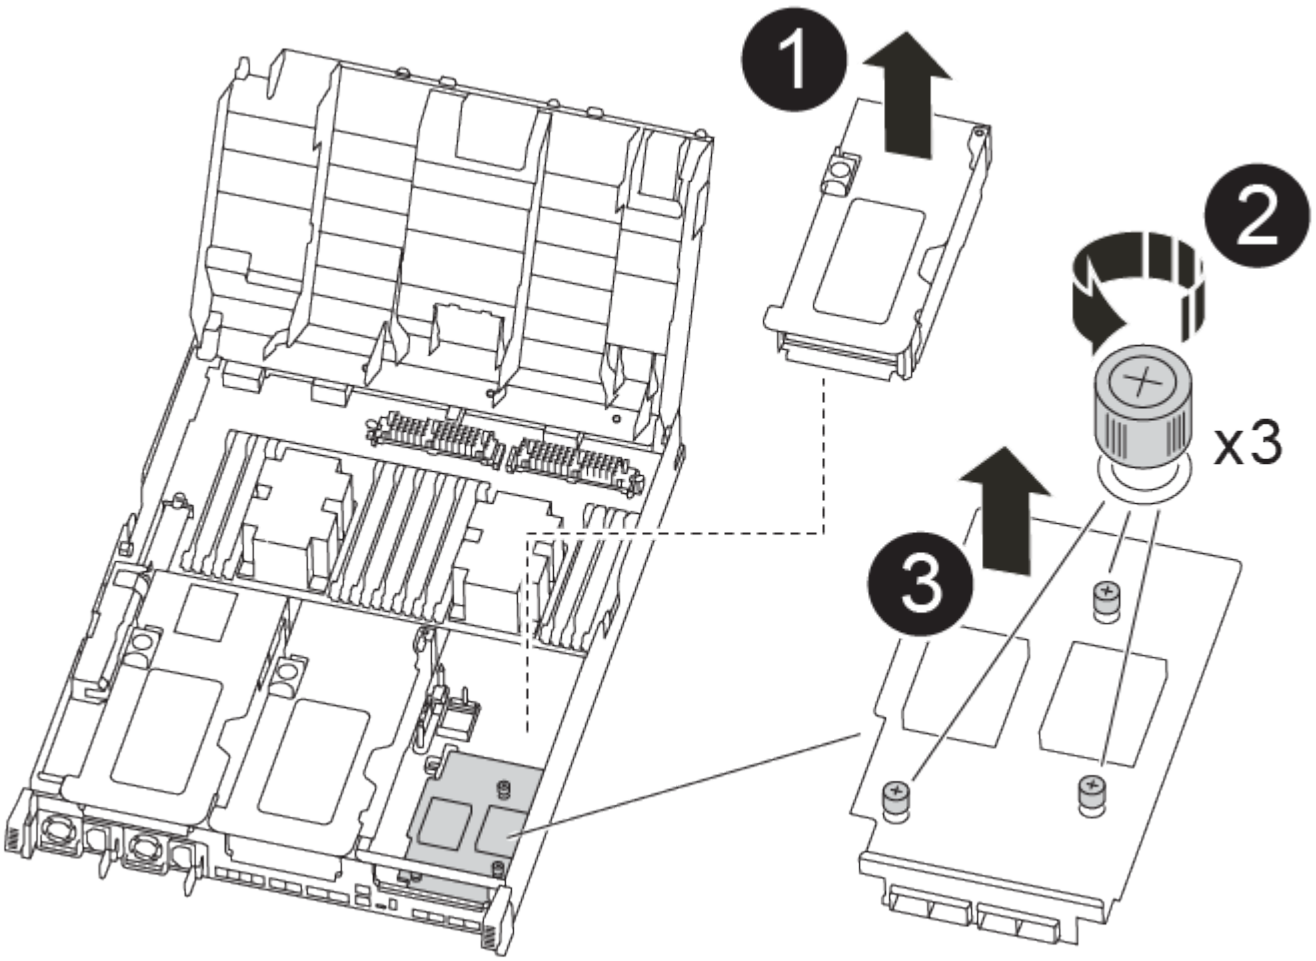

Steps

1. Remove riser number 3 (slots 4 and 5):
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up, and then set it aside on a stable, flat surface.
2. Replace the mezzanine card:
   a. Remove any QSFP or SFP modules from the card.
   b. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket and set it aside.
   c. Align the replacement mezzanine card over the socket and the guide pins and gently push the card into the socket.
   d. Tighten the thumbscrews on the mezzanine card.
3. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 5: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module

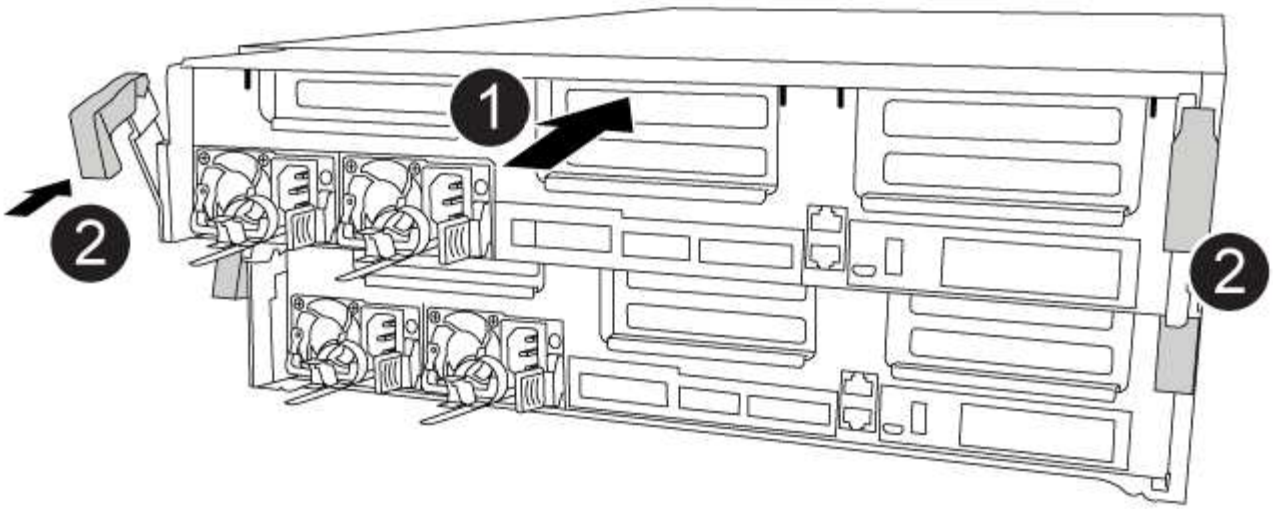

Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

5. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
6. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 8: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - FAS8300 and FAS8700

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting the replacement PSU to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

You can use the following animation, illustration, or the written steps to replace the power supply.

Animation - Replace a power supply

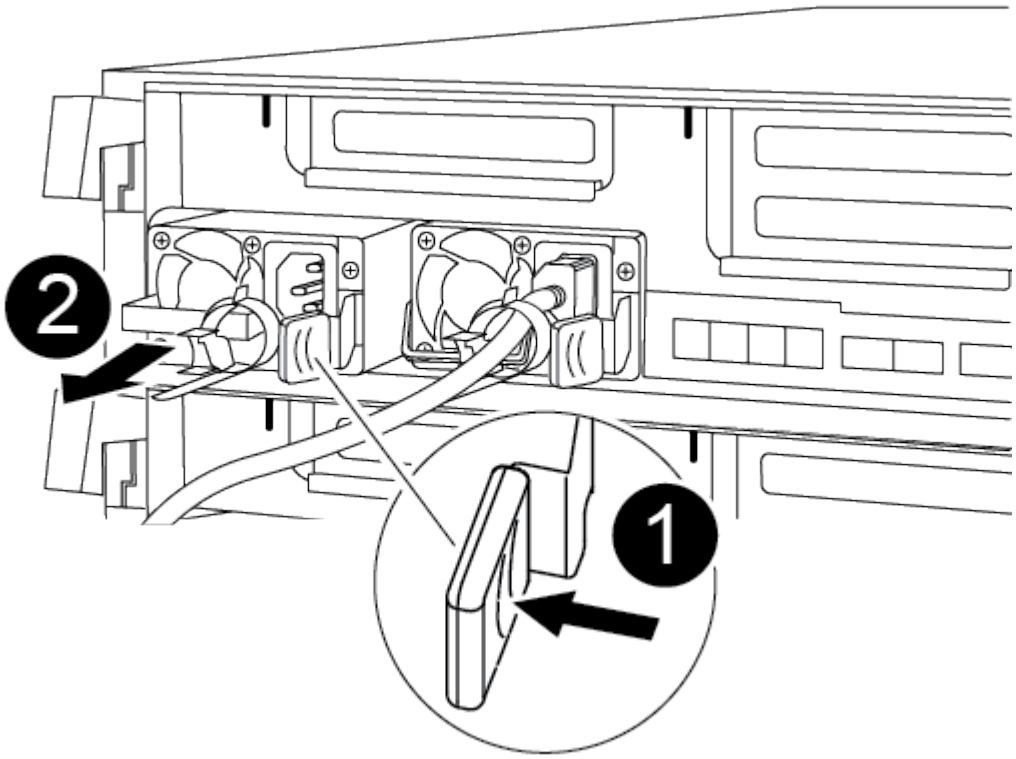

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
3. Disconnect the power supply:
   a. Open the power cable retainer, and then unplug the power cable from the power supply.
   b. Unplug the power cable from the power source.
4. Remove the power supply:
   a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
   b. Press the blue locking tab to release the power supply from the chassis.
   c. Using both hands, pull the power supply out of the chassis, and then set it aside.
5. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

6. Rotate the cam handle so that it is flush against the power supply.

7. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply and the power source.

   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.

8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - FAS8300 and FAS8700

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

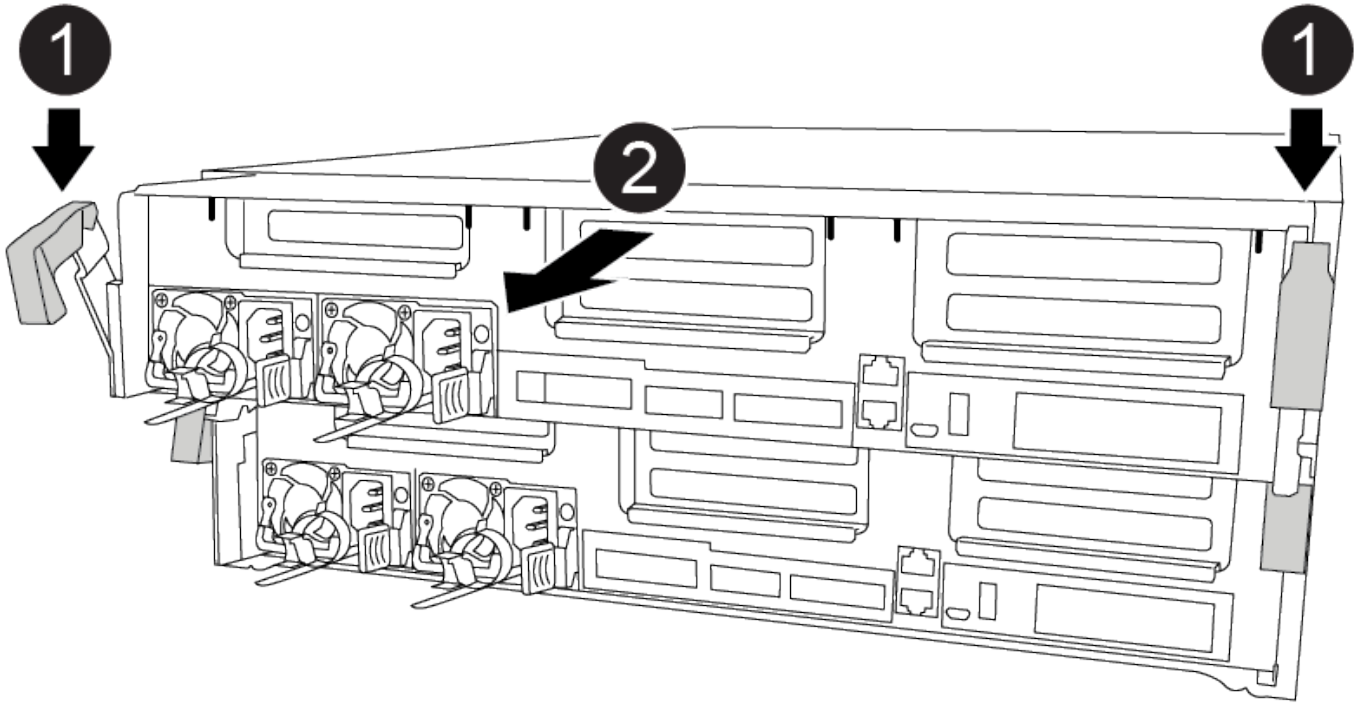

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Place the controller module on a stable, flat surface.

Step 3: Replace the RTC battery

You need to locate the RTC battery inside the controller module, and then follow the specific sequence of steps. See the FRU map inside the controller module for the location of the RTC battery.

You can use the following animation, illustration, or the written steps to replace the RTC battery.

Animation - Replace the RTC battery

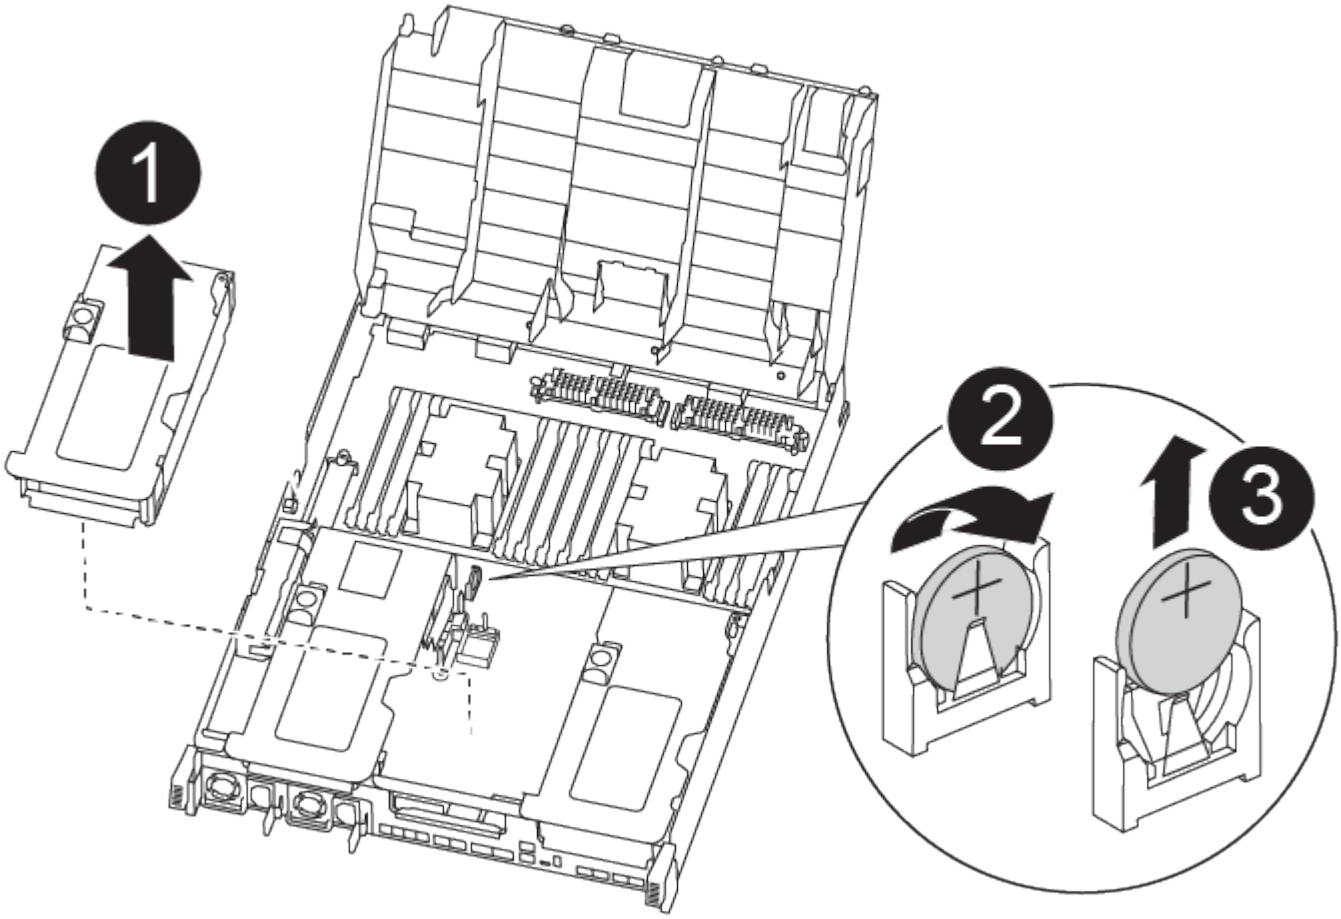

Steps

1. If you are not already grounded, properly ground yourself.
2. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
3. Locate, remove, and then replace the RTC battery:
   a. Using the FRU map, locate the RTC battery on the controller module.
   b. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

c. Remove the replacement battery from the antistatic shipping bag.

d. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.

4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

5. Close the air duct.

Step 4: Reinstall the controller module and sett time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module

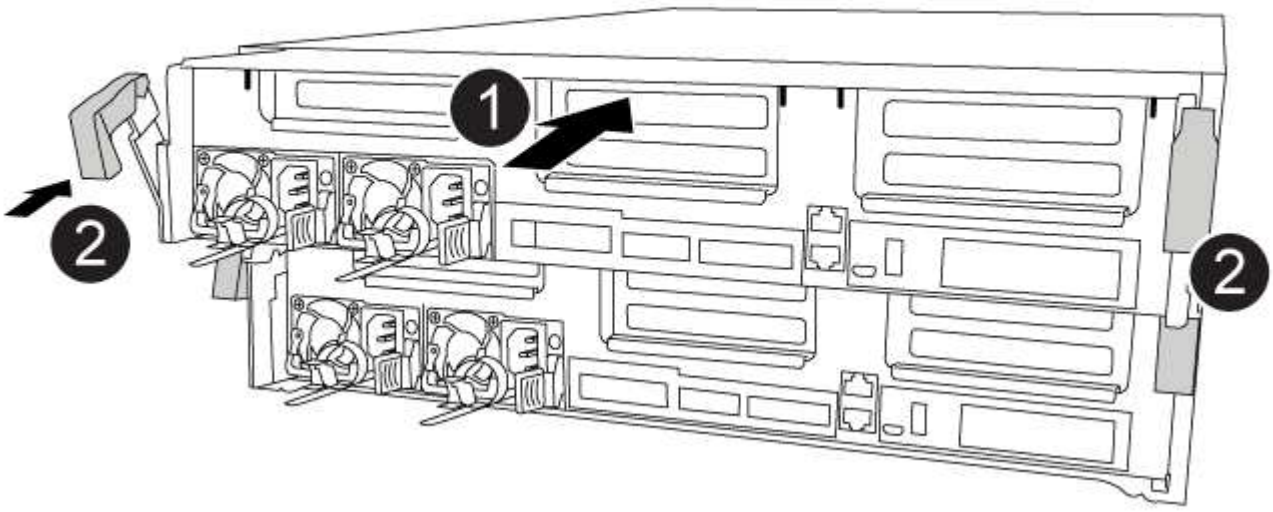

Steps

1. If you have not already done so, close the air duct or controller module cover.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      > If your system stops at the boot menu, select the option to boot to LOADER.

5. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

FAS9500 systems

Install and setup

Maintain

Maintain FAS9500 hardware

Maintain the hardware of your FAS9500 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS9500 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS9500 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Caching module	You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

DCPM	The DCPM (destage controller power module) contains the NVRAM11 battery.
Fan	The fan cools the controller.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
LED USB	The LED USB module provides connectivity to console ports and system status.
NVRAM	The NVRAM module (Non-Volatile Random Access Memory) allows the controller to retain data across power cycles or system reboots, while the NVRAM DIMM maintains NVRAM settings.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - FAS9500

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your FAS9500 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

Review the boot media requirements

Review the requirements for boot media replacement.

Shut down the controller

Shut down the controller in your storage system when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the controller module and install the replacement boot media.

4 Restore the image on the boot media

Restore the ONTAP image from the partner controller.

5 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - FAS9500

Before replacing the boot media in your FAS9500, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - FAS9500

Shut down the impaired controller in your FAS9500 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - FAS9500

The boot media in your FAS9500 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

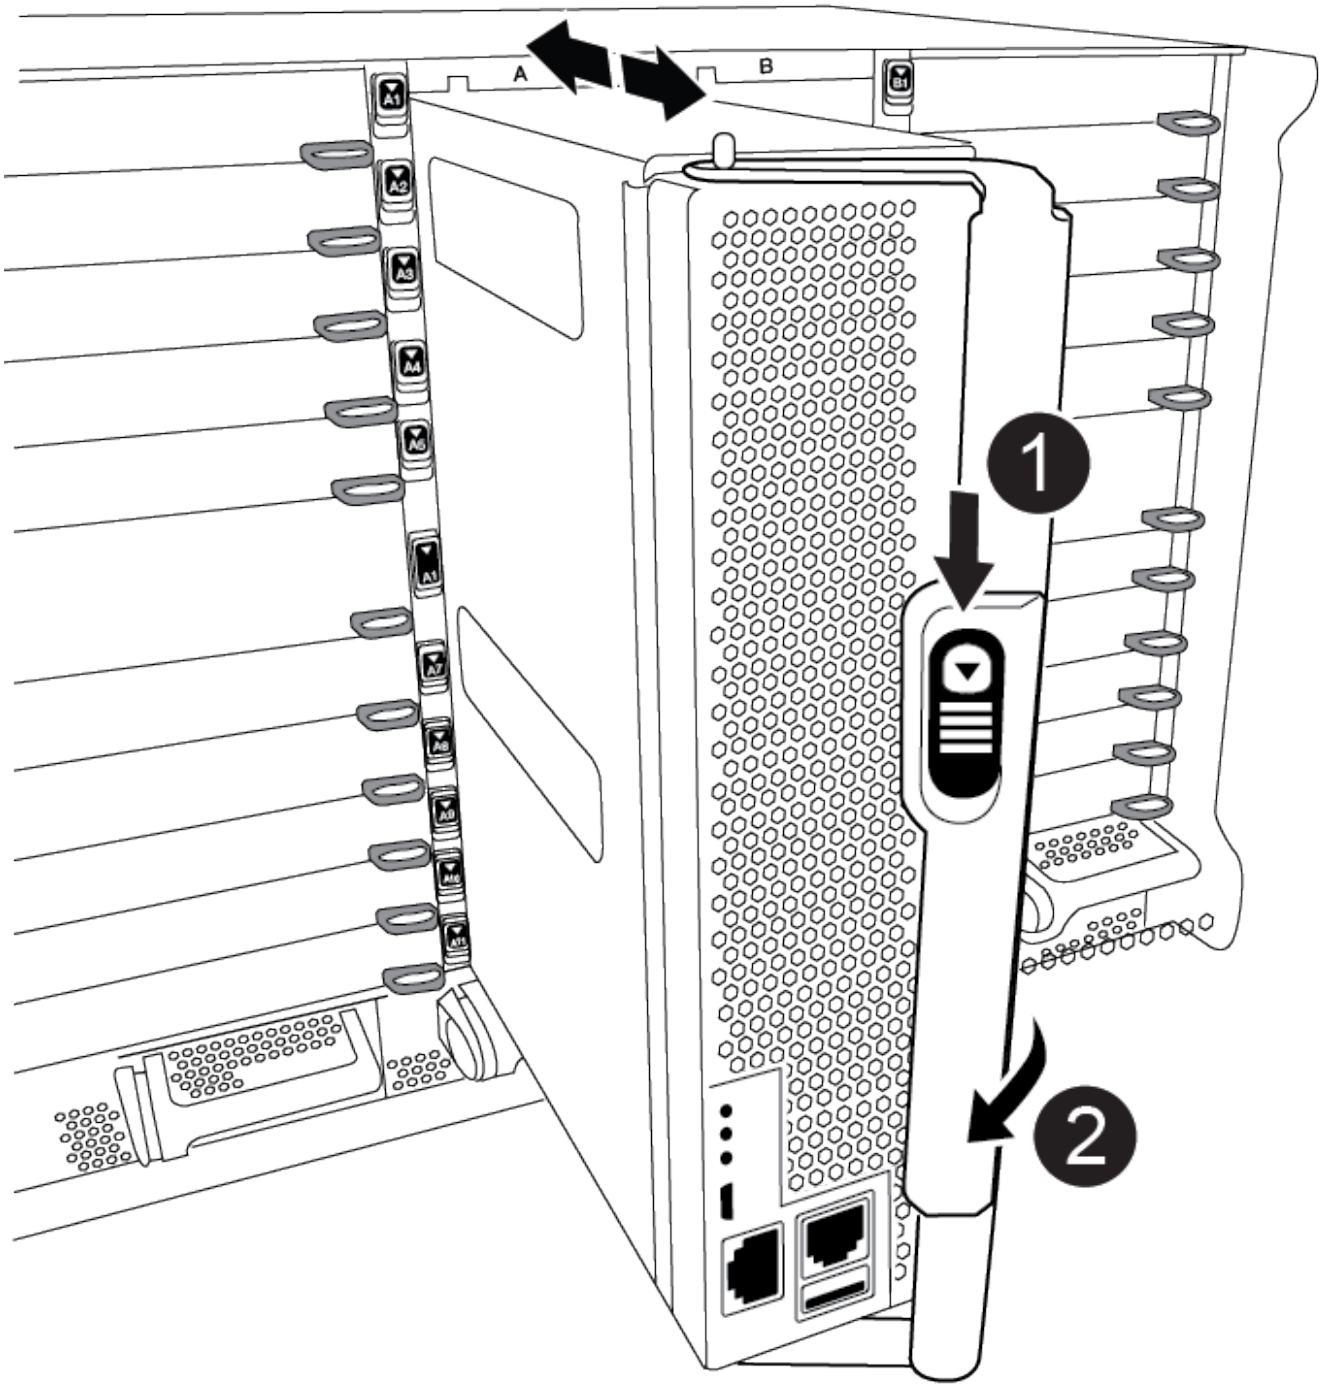

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

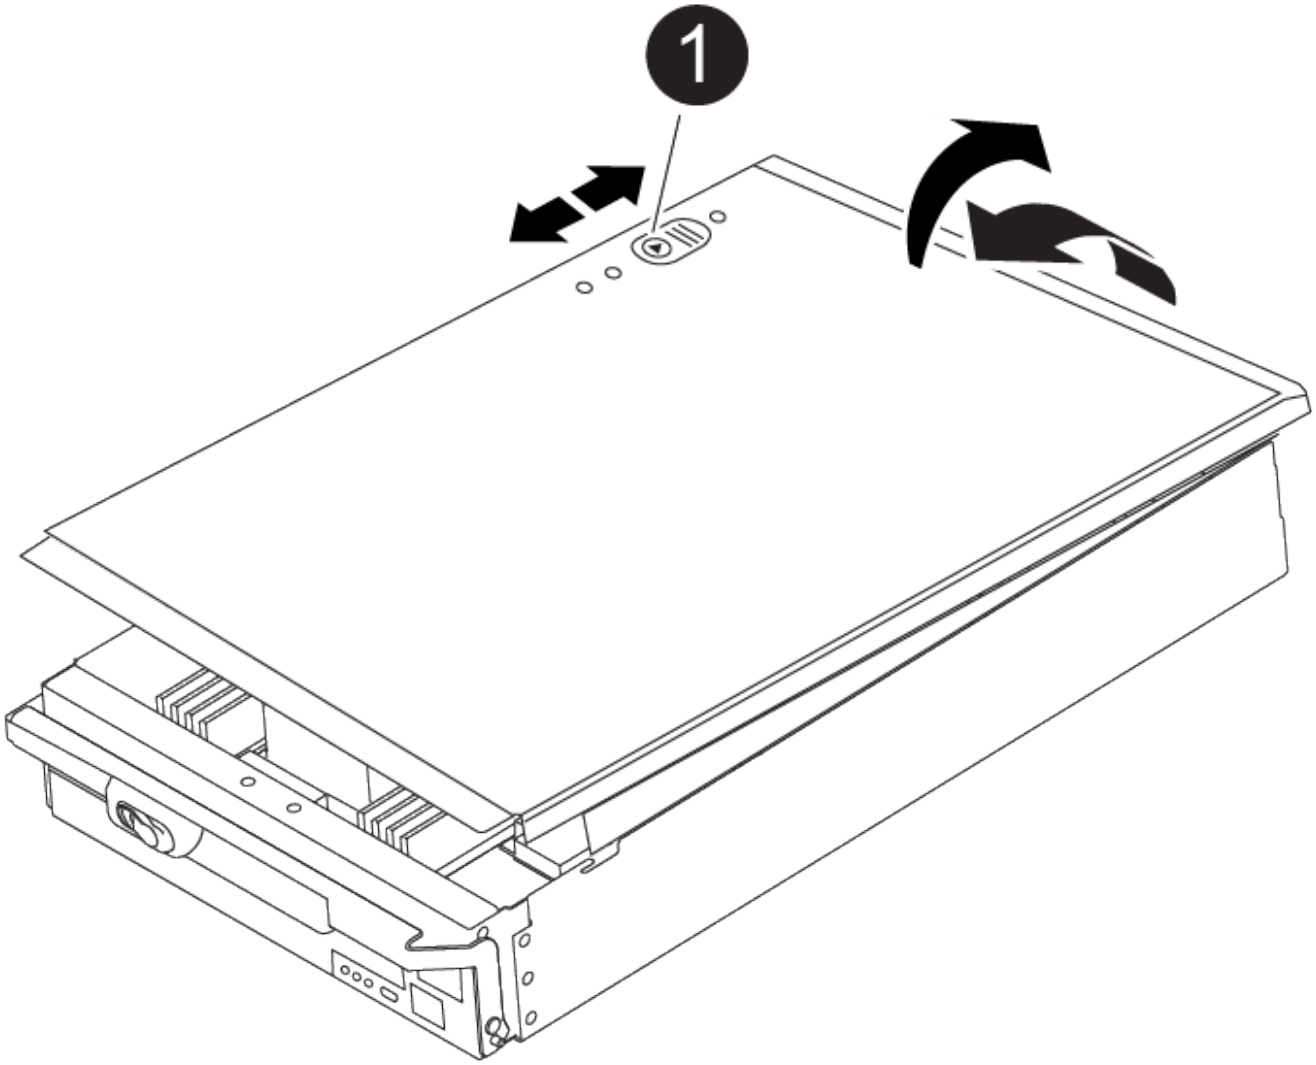

1	Controller module cover locking button

6. Replace the boot media:
   a. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

      Animation - Replace boot media

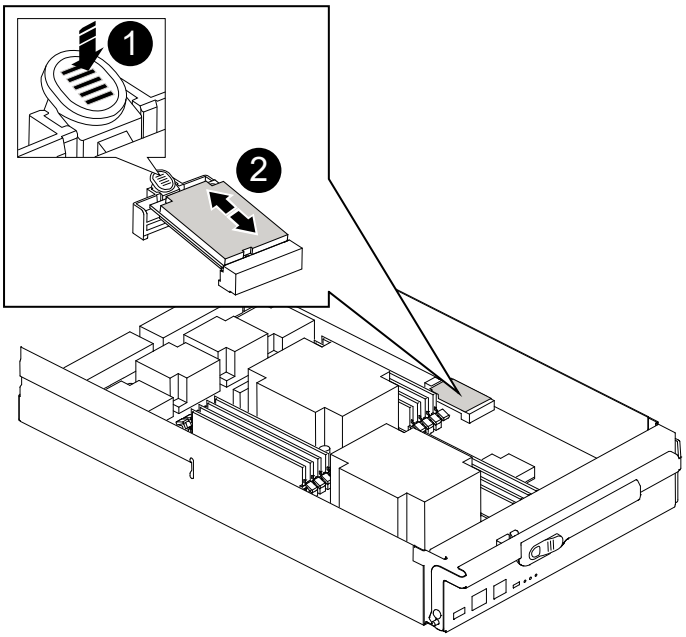

1	Press release tab
2	Boot media

b. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

c. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.

d. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

e. Push the boot media down to engage the locking button on the boot media housing.

7. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.
8. Reinstall the controller module:
   a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
   b. Recable the controller module, as needed.
   c. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, and then push the cam handle to the closed position.

      The controller begins to boot as soon as it is completely installed into the chassis.

If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

9. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t iniitator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - FAS9500

After installing the new boot media device in your FAS9500 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.

Steps

1. From the LOADER prompt, enter the command:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and encryption type and displays one of two messages. Depending on what message is displayed, take one of the following actions:

Occasionally, the process may not be able to identify if key manager is configured on the system. It will display an error message, ask if key manager is configured for the system, and then ask what type of key manager is configured. The process will resume after you resolve the issue.

Show example of configuration error finding prompts

```
Error when fetching key manager config from partner ${partner_ip}:
${status}

Has key manager been configured on this system

Is the key manager onboard
```

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. Complete the following steps: a. Log into the node when the login prompt is displayed and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* b. Go to step 5 to enable automatic giveback if it was disabled.
`key manager is configured.`	Go to step 4 to restore the appropriate key manager. The node accesses the boot menu and runs: • Option 10 for systems with Onboard Key Manager (OKM). • Option 11 for systems with External Key Manager (EKM).

4. Select the appropriate key manager restoration process.

Onboard Key Manager (OKM)

If OKM is detected, the system displays the following message and begins running BootMenu Option 10.

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `Y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the following when prompted:

   i. The passphrase

   ii. The passphrase again when prompted to confirm

   iii. Backup data for onboard key manager

   Show example of passphrase and backup data prompts

   ```
   Enter the passphrase for onboard key management:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the passphrase again to confirm:
   -----BEGIN PASSPHRASE-----
   <passphrase_value>
   -----END PASSPHRASE-----
   Enter the backup data:
   -----BEGIN BACKUP-----
   <passphrase_value>
   -----END BACKUP-----
   ```

c. Continue to monitor the recovery process as it restores the appropriate files from the partner node.

   When the recovery process is complete, the node will reboot. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

d. When the node reboots, verify the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

f. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster.

```
security key-manager onboard sync
```

External Key Manager (EKM)

If EKM is detected, the system displays the following message and begins running BootMenu Option 11.

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. The next step depends on which version of ONTAP your system is running:

If your system is running…	Do this…
ONTAP 9.16.0	a. Press `Ctlr-C` to exit BootMenu Option 11. b. Press `Ctlr-C` to exit the EKM configuration process and return to the boot menu. c. Select BootMenu Option 8. d. Reboot the node. If `AUTOBOOT` is set, the node reboots and uses the configuration files from the partner node. If `AUTOBOOT` is not set, enter the appropriate boot command. The node reboots and uses the configuration files from the partner node. e. Reboot the node so that EKM protects the boot media partition. f. Proceed to step c.

If your system is running…	Do this…
ONTAP 9.16.1 and later	Proceed to the next step.

b. Enter the following EKM configuration setting when prompted:

Action	Example

Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file.	**Show example of client certificate contents**
	`-----BEGIN CERTIFICATE-----`
	`<certificate_value>`
	`-----END CERTIFICATE-----`

Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file.	**Show example of client key file contents**
	`-----BEGIN RSA PRIVATE KEY-----`
	`<key_value>`
	`-----END RSA PRIVATE KEY-----`

Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file.	**Show example of KMIP server file contents**
	`-----BEGIN CERTIFICATE-----`
	`<KMIP_certificate_CA_value>`
	`-----END CERTIFICATE-----`

Action	Example

Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file.	**Show example of server configuration file contents**
	`xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx`
	`xxx.xxx.xxx.xxx:5696.port=5696`
	`xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem`
	`xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4`
	`1xxx.xxx.xxx.xxx:5696.timeout=25`
	`xxx.xxx.xxx.xxx:5696.nbio=1`
	`xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt`
	`xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key`
	`xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:!RC2:!RC4:!SEED:!eNULL:!aNULL"`
	`xxx.xxx.xxx.xxx:5696.verify=true`
	`xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>`

Action	Example
If prompted, enter the ONTAP Cluster UUID from the partner. You can check the cluster UUID from the partner node using the `cluster identify show` command.	**Show example of ONTAP Cluster UUID** ``` Notice: bootarg.mgwd.cluster_uuid is not set or is empty. Do you know the ONTAP Cluster UUID? {y/n} y Enter the ONTAP Cluster UUID: <cluster_uuid_value> System is ready to utilize external key manager(s). ```
If prompted, enter the temporary network interface and settings for the node. You need to enter: 1. The IP address for the port 2. The netmask for the port 3. The IP address of the default gateway	**Show example of a temporary network setting** ``` In order to recover key information, a temporary network interface needs to be configured. Select the network port you want to use (for example, 'e0a') e0M Enter the IP address for port : xxx.xxx.xxx.xxx Enter the netmask for port : xxx.xxx.xxx.xxx Enter IP address of default gateway: xxx.xxx.xxx.xxx Trying to recover keys from key servers.... [discover_versions] [status=SUCCESS reason= message=] ```

c. Depending on whether the key is successfully restored, take one of the following actions:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored.

  The process attempts to restore the appropriate files from the partner node and reboots the node. Go to step d.

- If the key is not successfully restored, the system will halt and indicate that it could not restore the key. The error and warning messages are displayed. You must rerun the recovery process:

  ```
  boot_recovery -partner
  ```

  Show example of key recovery error and warning messages

  ```
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  WARNING: kmip_init: authentication keys might not be
  available.
  ********************************************************
  *                A T T E N T I O N                     *
  *                                                      *
  *       System cannot connect to key managers.         *
  *                                                      *
  ********************************************************
  ERROR: kmip_init: halting this system with encrypted
  mroot...
  .
  Terminated

  Uptime: 11m32s
  System halting...

  LOADER-B>
  ```

d. When the node reboots, verify that the boot media recovery was successful by confirming that the system is back online and operational.

e. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - FAS9500

If a component in your FAS9500 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - FAS9500

Get started with replacing the boot media in your FAS9500 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Review the boot media requirements

Review the requirements for replacing the boot media.

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

Shut down the controller

Shut down the controller when when you need to replace the boot media.

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6 Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.

7 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - FAS9500

Before replacing the boot media in your FAS9500 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - FAS9500

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller for manual boot media recovery - FAS9500

Shut down or take over the impaired controller using one of the following options.

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the boot media and prepare for manual boot recovery - FAS9500

You must unplug the controller module, remove and open the controller module, locate and replace the boot media in the controller, and then transfer the image to the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and

then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

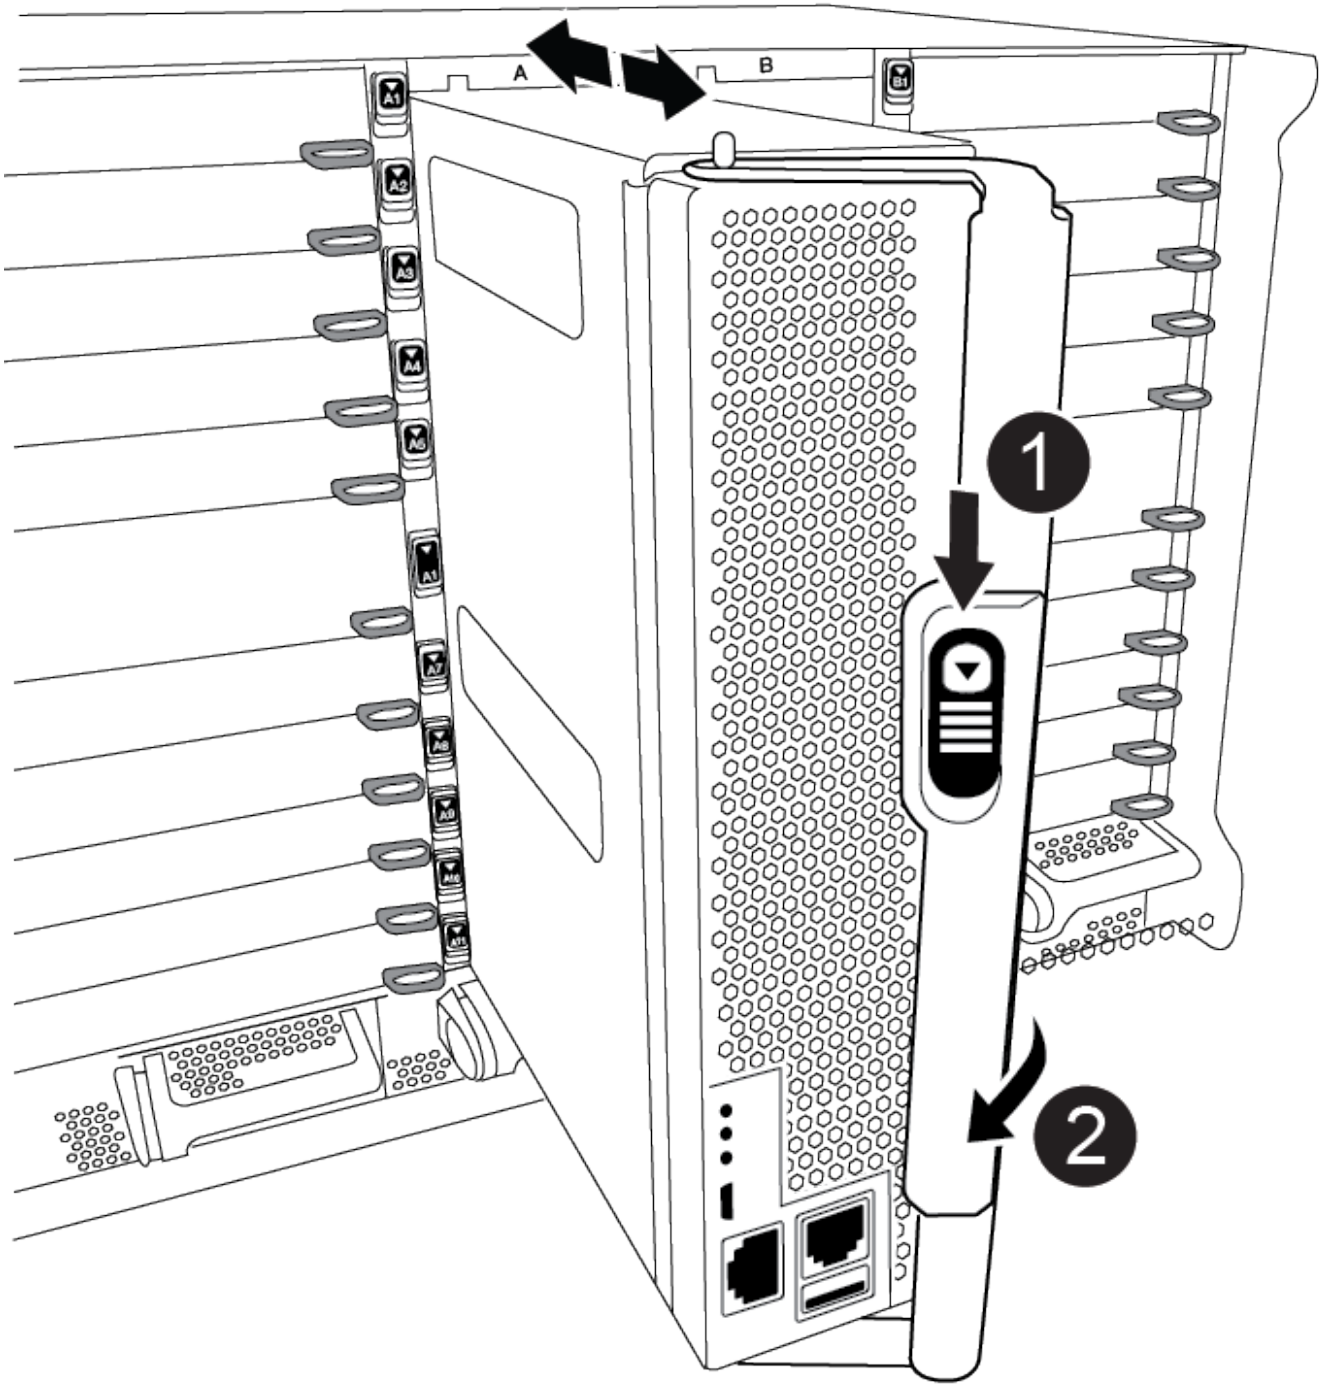

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

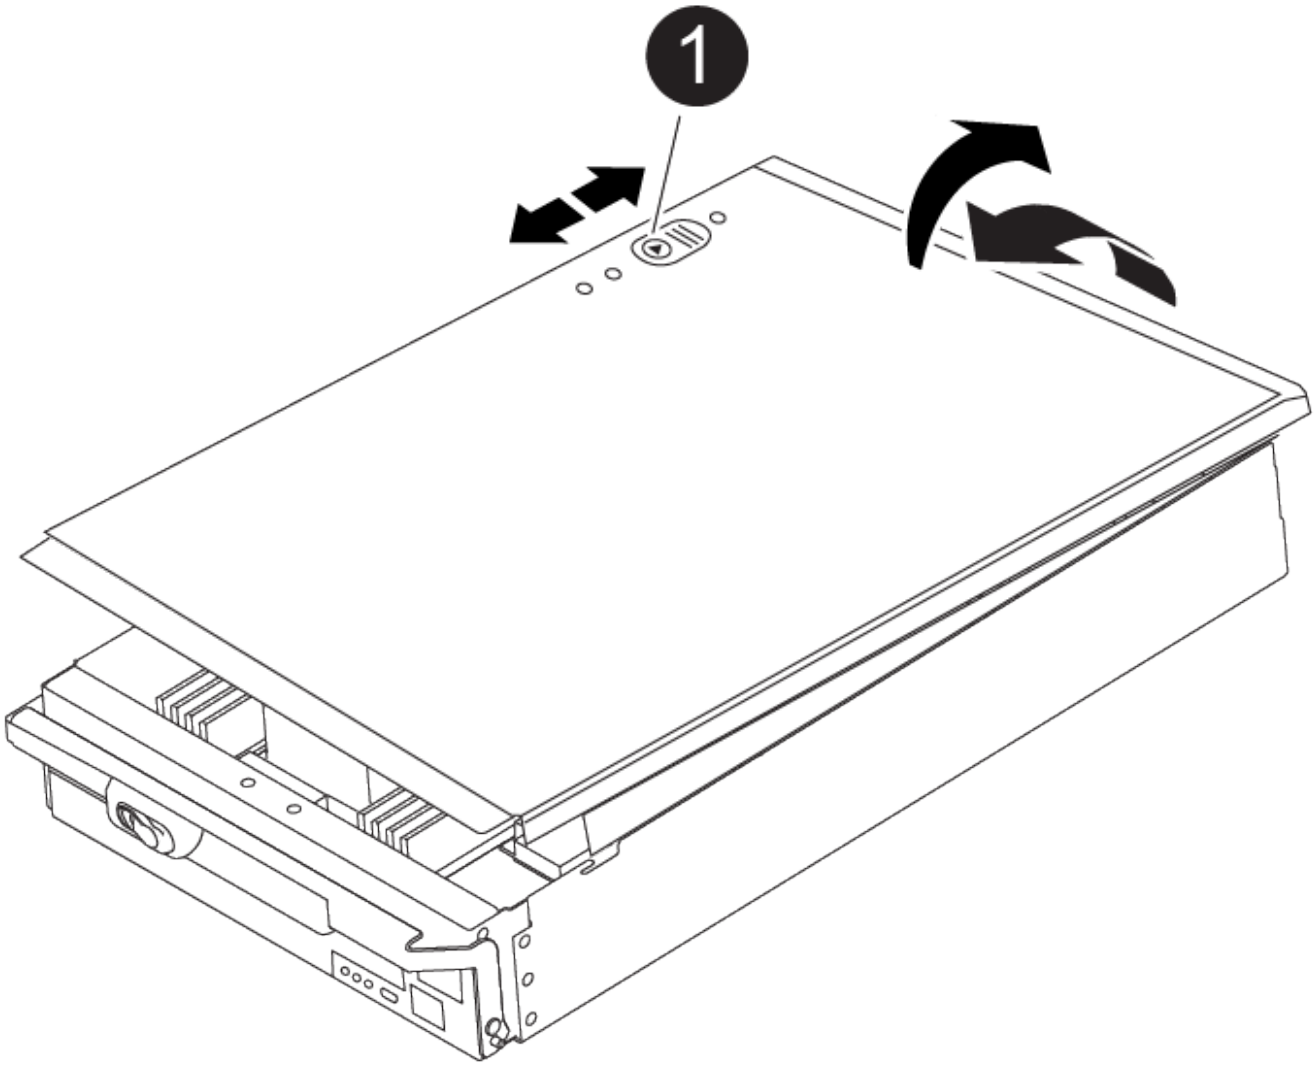

1	Controller module cover locking button

Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

   Animation - Replace boot media

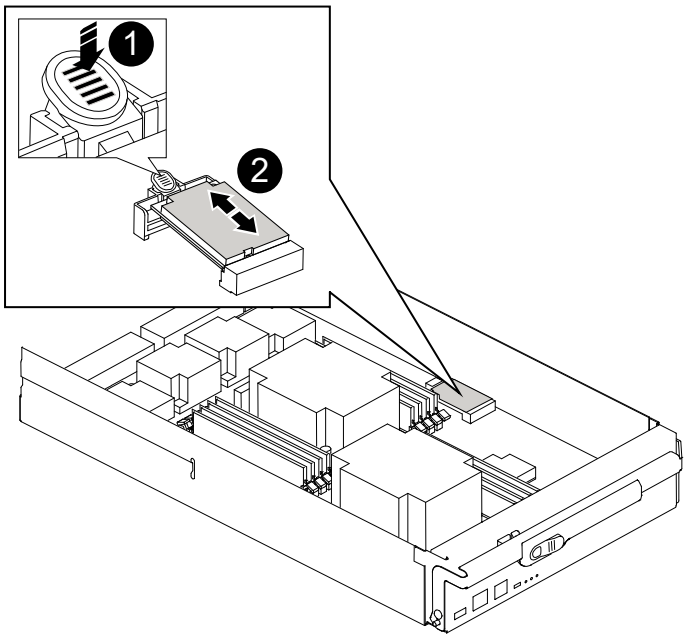

1	Press release tab
2	Boot media

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.
6. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- Download a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site. Use the `version -v` command to display if your version of ONTAP supports NVE. If the command output displays `<10no- DARE>`, your version of ONTAP does not support NVE.
  - If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. If you have not done so, download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image from the Downloads link on the page, to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      The USB flash drive should have the appropriate ONTAP image of what the impaired controller is running.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
3. Recable the controller module, as needed.
4. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.
5. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, and then push the cam handle to the closed position.

   The controller begins to boot as soon as it is completely installed into the chassis.
6. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

7. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:

   a. Boot to Maintenance mode: `boot_ontap maint`

   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`

   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - FAS9500

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

Option 1: ONTAP 9.16.0 or earlier

a. On the impaired controller, press `Y` when you see `Do you want to restore the backup configuration now?`

b. If prompted on the impaired controller, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. On the healthy partner controller, set the impaired controller to advanced privilege level: `set -privilege advanced`.

d. On the healthy partner controller, run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address`.

NOTE: If you see any message other than a successful restore, contact NetApp Support.

e. On the healthy partner controller, return the impaired controller to admin level: `set -privilege admin`.

f. On the impaired controller, press `Y` when you see `Was the restore backup procedure successful?`.

g. On the impaired controller, press `Y` when you see `…would you like to use this restored copy now?`.

h. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

i. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

Option 2: ONTAP 9.16.1 or later

a. On the impaired controller, press `Y` when prompted to restore the backup configuration.

After restore procedure is successful, this message will be seen on the console - `syncflash_partner: Restore from partner complete`.

b. On the impaired controller, press `Y` when prompted to confirm if the restore backup was successful.

c. On the impaired controller, press `Y` when prompted to use the restored configuration.

d. On the impaired controller, press `Y` when prompted to reboot the node.

e. On the impaired controller, press `Y` when prompted to reboot the impaired controller and press `ctrl-c` for the Boot Menu.

f. If the system does not use encryption, select *Option 1 Normal Boot.*, otherwise go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Give back the controller using the `storage failover giveback -fromnode local` command.
6. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
7. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

NOTE: If the process fails, contact NetApp Support.

Restore encryption - FAS9500

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*******************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*******************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - FAS9500

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Replace the chassis - FAS9500

Before you begin

To replace the chassis, you must remove the power supplies, fans, controller modules, I/O modules, DCPM modules, and USB LED module from the impaired chassis, remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis in its place, and then install the components into the replacement chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

About this task

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shutdown the impaired controller - FAS9500

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - FAS9500

To replace the chassis, you must remove the components from the impaired chassis and install them in the replacement chassis.

Step 1: Remove the power supplies

Removing the power supplies when replacing a chassis involves turning off, disconnecting, and then removing the four power supplies from the rear of the impaired chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Press and hold the terra cotta locking button on the power supply handle, and then pull the power supply out of the chassis.

> When removing a power supply, always use two hands to support its weight.

Animation - Remove/install PSU

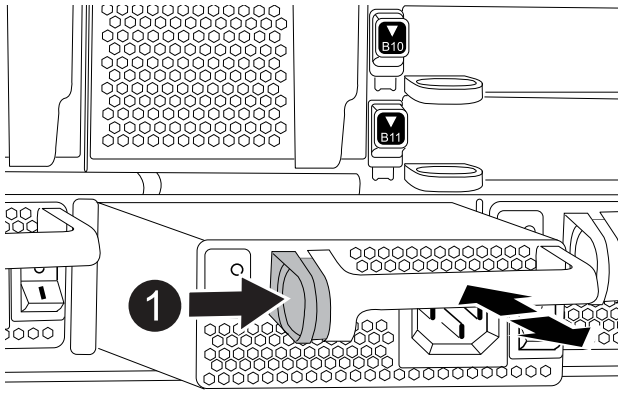

1	Terra cotta locking button

4. Repeat the preceding steps for any remaining power supplies.

Step 2: Remove the fans

You must remove the six fan modules, located on in the front of the chassis, when replacing the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press the terra cotta locking button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

> The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

Animation - Remove/install fan

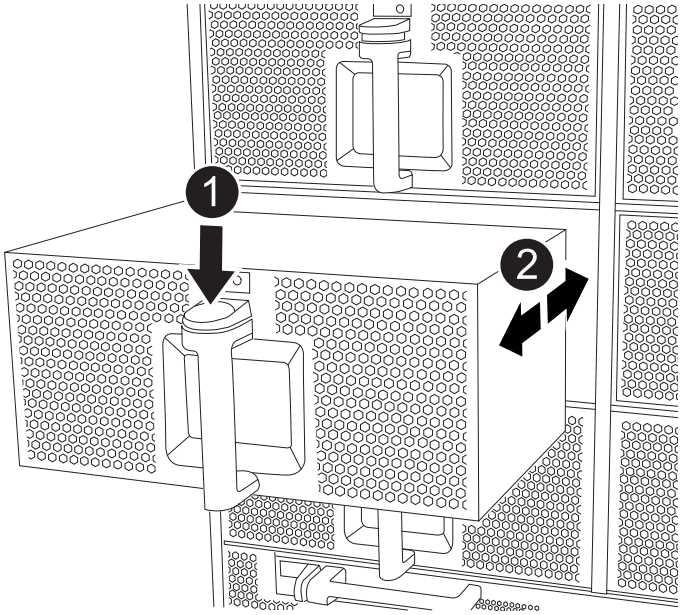

1	Terra cotta locking button
2	Slide fan in/out of chassis

4. Set the fan module aside.
5. Repeat the preceding steps for any remaining fan modules.

Step 3: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the impaired chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta locking button on the cam handle downward until it unlocks.

   Animation - Remove controller module

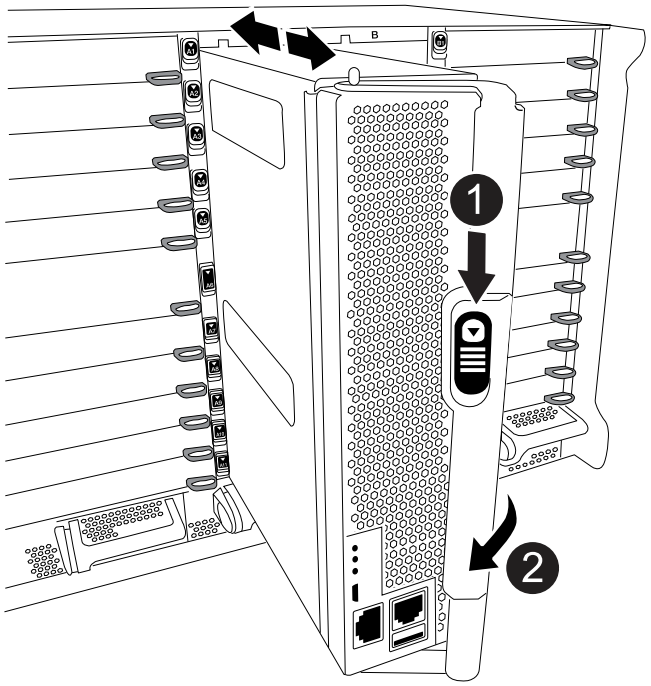

1	Cam handle locking button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Set the controller module aside in a safe place and keep track of which chassis slot it came from, so that it can be installed into the same slot in the replacement chassis..
6. Repeat these steps if you have another controller module in the chassis.

Step 4: Remove the I/O modules

To remove I/O modules from the impaired chassis, including the NVRAM modules, follow the specific sequence of steps. You do not have to remove the Flash Cache module, if present, from the NVRAM module

when moving it to a replacement chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

3. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam locking button.

      The cam locking button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

      Animation - Remove/install I/O module

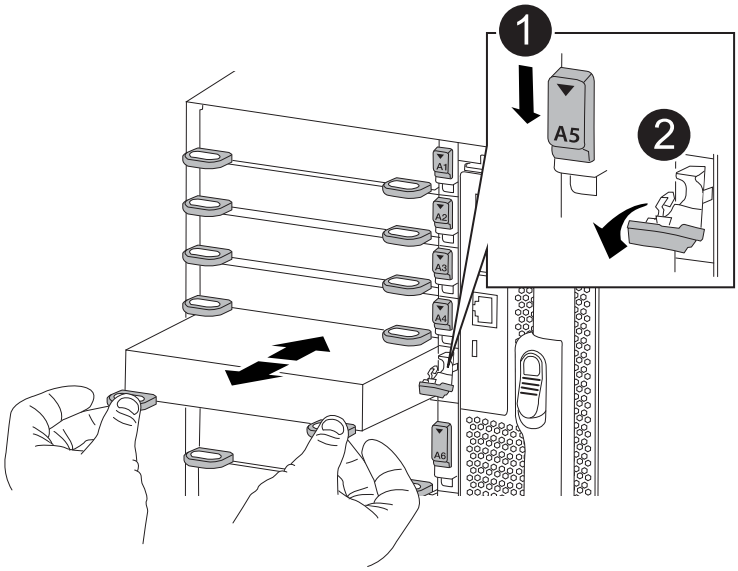

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Repeat the preceding step for the remaining I/O modules in the impaired chassis.

Step 5: Remove the De-stage Controller Power Module

Remove the two de-stage controller power modules from the front of the impaired chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Press the terra cotta locking button on the module handle, and then slide the DCPM out of the chassis.

   Animation - Remove/install DCPM

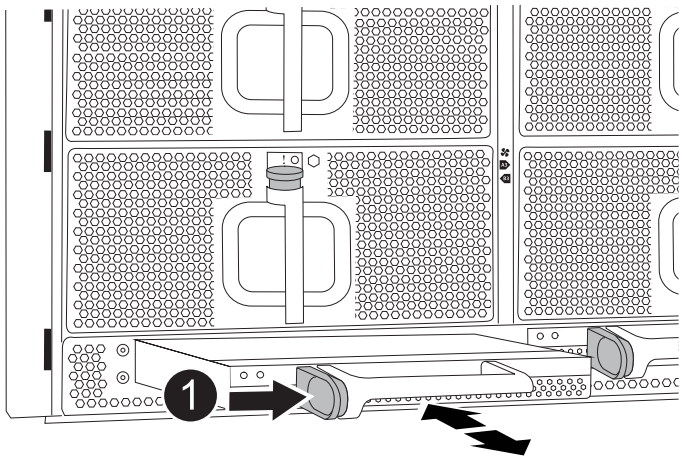

1	DCPM terra cotta locking button

3. Set the DCPM aside in a safe place and repeat this step for the remaining DCPM.

Step 6: Remove the USB LED module

Remove the USB LED modules.

Animation - Remove/install USB module

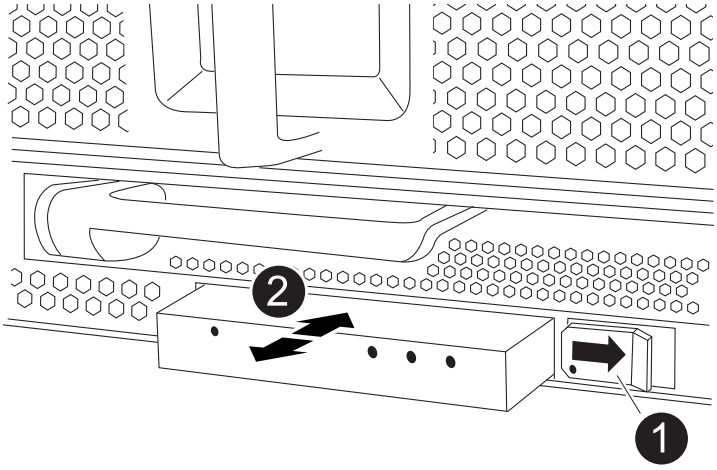

1	Eject the module.
2	Slide out of chassis.

Steps

1. Locate the USB LED module on the front of the impaired chassis, directly under the power supply bays.
2. Press the black locking button on the right side of the module to release the module from the chassis, and then slide it out of the impaired chassis.
3. Set the module aside in a safe place.

Step 7: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

Steps

1. Remove the screws from the chassis mount points.

   If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

2. With the help of two or three people, slide the impaired chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.
7. Secure the rear of the chassis to the equipment rack or system cabinet.
8. If you are using the cable management brackets, remove them from the impaired chassis, and then install them on the replacement chassis.

Step 8: Install the de-stage controller power module when replacing the chassis

Once the replacement chassis is installed into the rack or system cabinet, you must reinstall the de-stage controller power modules into it.

Steps

1. If you are not already grounded, properly ground yourself.
2. Align the end of the DCPM with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

3. Repeat this step for the remaining DCPM.

Step 9: Install fans into the chassis

To install the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

Steps

1. If you are not already grounded, properly ground yourself.
2. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

3. Repeat these steps for the remaining fan modules.
4. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Step 10: Install I/O modules

To install I/O modules, including the NVRAM/Flash Cache modules from the impaired chassis, follow the specific sequence of steps.

You must have the chassis installed so that you can install the I/O modules into the corresponding slots in the replacement chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. After the replacement chassis is installed in the rack or cabinet, install the I/O modules into their corresponding slots in the replacement chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage, and then push the I/O cam latch all the way up to lock the module in place.
3. Recable the I/O module, as needed.
4. Repeat the preceding step for the remaining I/O modules that you set aside.

   If the impaired chassis has blank I/O panels, move them to the replacement chassis at this time.

Step 11: Install the power supplies

Installing the power supplies when replacing a chassis involves installing the power supplies into the replacement chassis, and connecting to the power source.

Steps

1. If you are not already grounded, properly ground yourself.
2. Make sure the power supplies rockers are in the off position.
3. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

4. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

5. Repeat the preceding steps for any remaining power supplies.

Step 12 Install the USB LED modules

Install the USB LED modules in the replacement chassis.

Steps

1. Locate the USB LED module slot on the front of the replacement chassis, directly under the DCPM bays.
2. Align the edges of the module with the USB LED bay, and gently push the module all the way into the chassis until it clicks into place.

Step 13: Install the controller

After you install the controller module and any other components into the replacement chassis, boot the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Connect the power supplies to different power sources, and then turn them on.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the console to the controller module, and then reconnect the management port.
5. With the cam handle in the open position, slide the controller module into the chassis and firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle until it clicks into the locked position.

   Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.

   The controller module begins to boot as soon as it is fully seated in the chassis.

6. Repeat the preceding steps to install the second controller into the replacement chassis.
7. Boot each controller.

Restore and verify the configuration - FAS9500

To complete the chassis replacement, you must complete specific tasks.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and

chassis: `ha-config show`

The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis ha-state`

      The value for HA-state can be one of the following:

      - ha
      - non-ha
3. Confirm that the setting has changed: `ha-config show`
4. If you have not already done so, recable the rest of your system.

Step 2: Bring up the system

1. If you have not done so, plug the power cables back into the PSUs.
2. Turn on the PSUs by toggling the rocker switched to **ON**, and wait for the controllers to power up completely.
3. Check the front and the back of the chassis and controllers for any fault lights after power up.
4. Connect to the SP or BMC IP address of the nodes via SSH. This will be the same address used to shut down the nodes.
5. Perform additional health checks as described in How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP
6. Turn AutoSupport back on (end the maintenance window message):
   `system node autosupport invoke -node * -type all -message MAINT=end`

As a best practice, you should do the following:

- Resolve any Active IQ Wellness Alerts and Risks (Active IQ will take time to process post-power up AutoSupports - expect a delay in results)
- Run Active IQ Config Advisor
- Check system health using How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Replace the controller module - FAS9500

To replace the impaired controller module, you must shut down the impaired controller, move the internal components to the replacement controller module, install the replacement controller module, and reboot the replacement controller.

Before you begin

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system has a V_StorageAttach license, you must refer to the additional required steps before performing this procedure.
- If your system is in an HA pair, the healthy node must be able to take over the node that is being replaced (referred to in this procedure as the “impaired node”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a node in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired node to the replacement node so that the replacement node will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The impaired node is the node that is being replaced.
  - The replacement node is the new node that is replacing the impaired node.
  - The healthy node is the surviving node.
- You must always capture the node’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired node - FAS9500

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - FAS9500

To replace the controller module hardware, you must remove the impaired node, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

The following animation shows the whole process of moving components from the impaired to the replacement controller.

Animation - Replace controller module, complete process

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove controller module

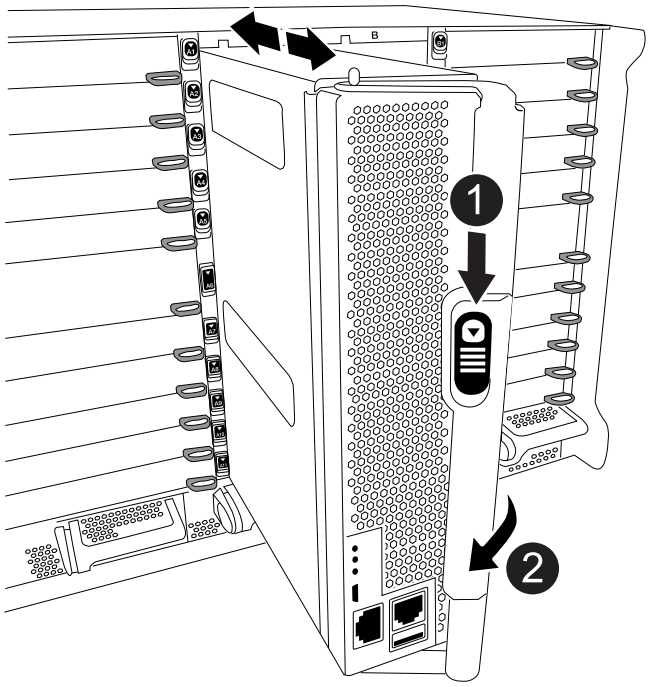

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

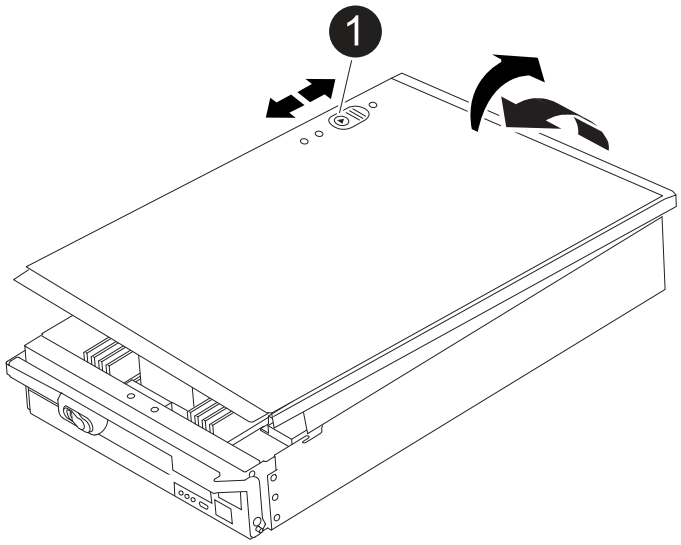

1	Controller module cover locking button

Step 2: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller and insert it in the new controller.

Steps

1. Locate the boot media using the following illustration or the FRU map on the controller module:

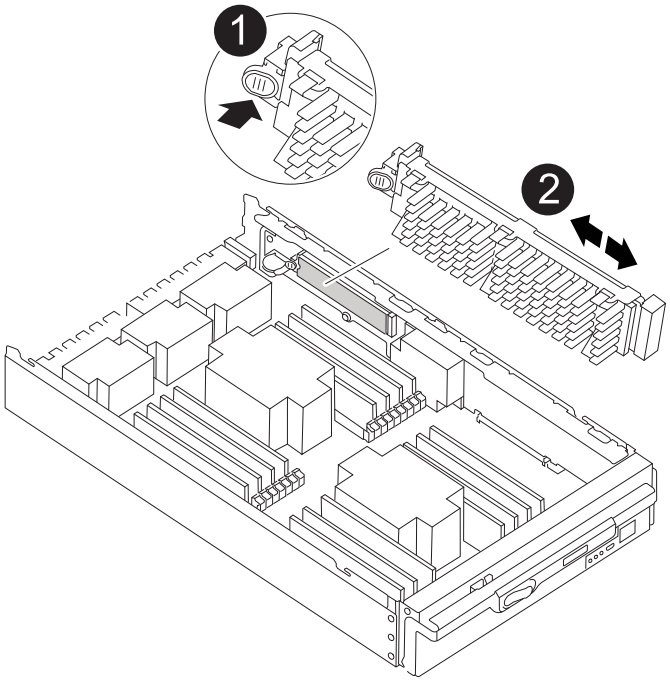

1	Press release tab

2	Boot media

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the system DIMMs

To move the DIMMs, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

The VER2 controller has fewer DIMM sockets. There is no reduction in the number of DIMMs supported or change in the DIMM socket numbering. When moving the DIMMs to the new controller module, install the DIMMs into the same socket number/location as the impaired controller module. See the FRU map diagram on the VER2 controller module for DIMM socket locations.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

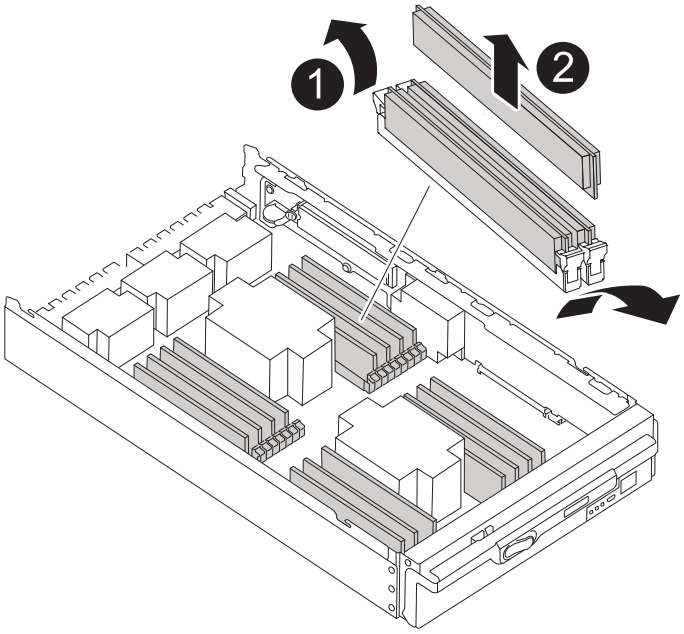

1	DIMM ejector tabs
2	DIMM

5. Locate the slot where you are installing the DIMM.
6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
9. Repeat these steps for the remaining DIMMs.

Step 4: Install the controller

After you install the components into the replacement controller module, you must install the replacement controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Animation - Install controller module

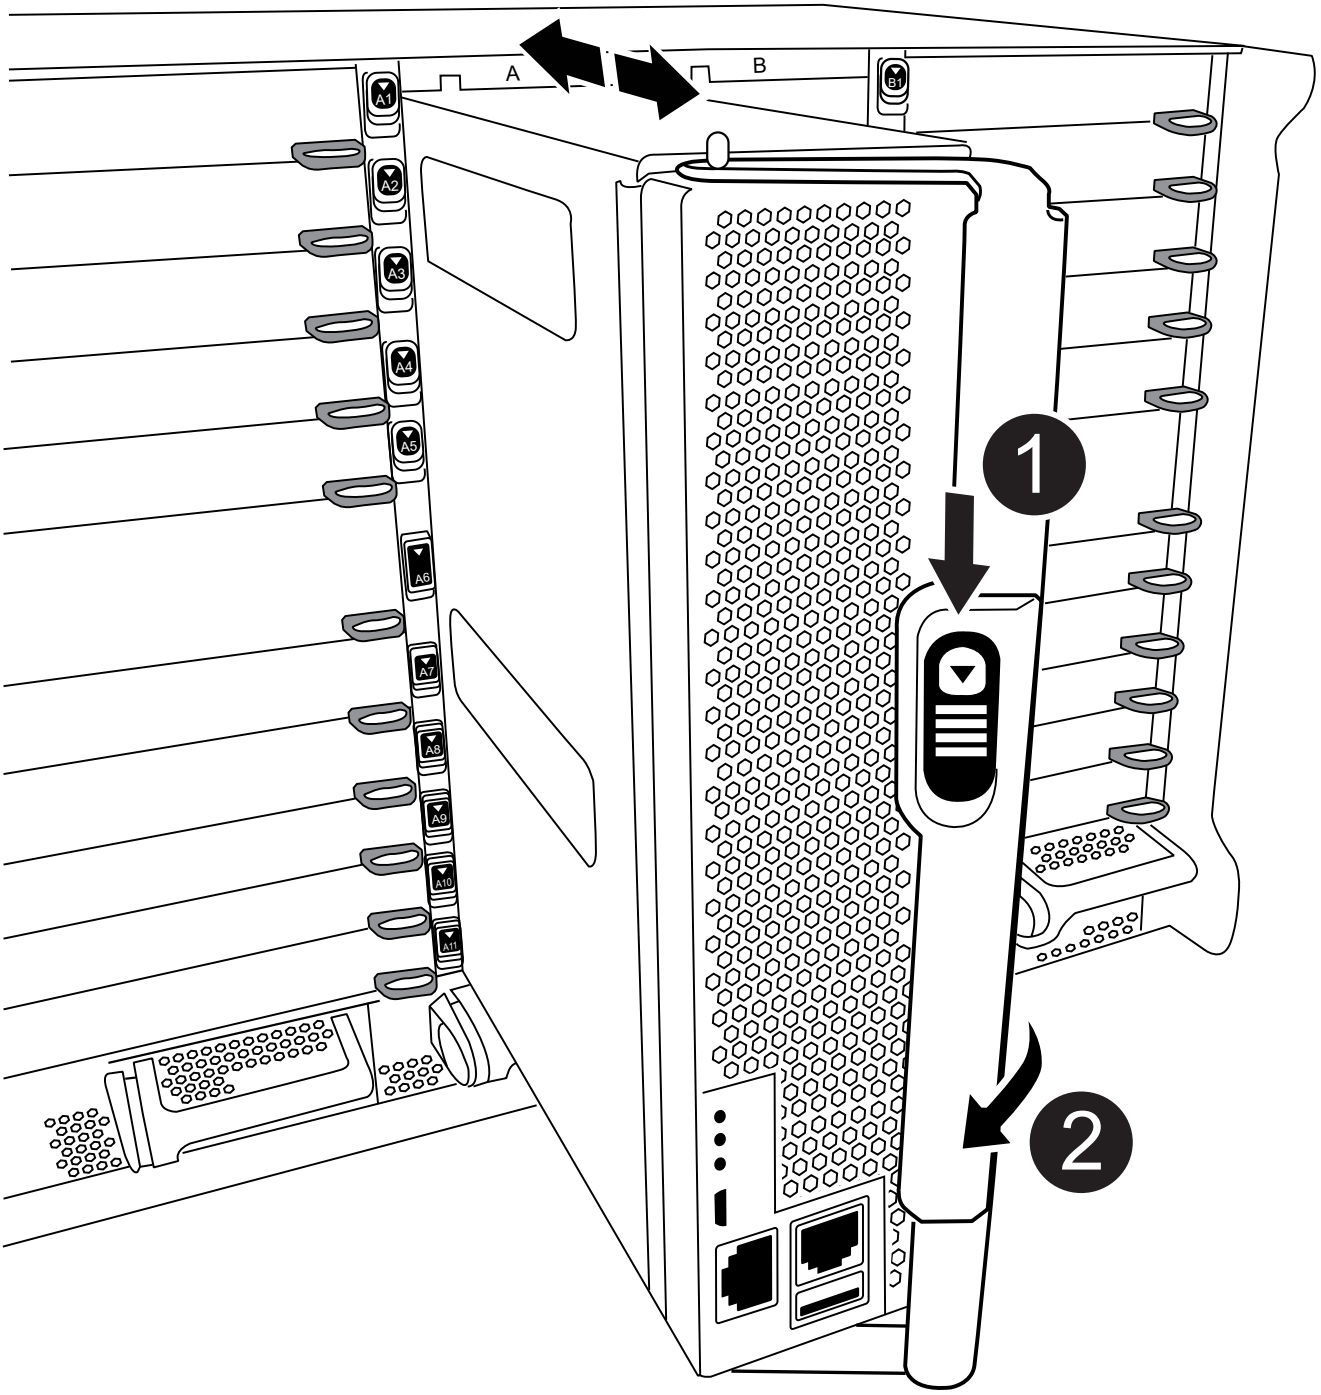

1	Cam handle release button
2	Cam handle

Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   c. Rotate the controller module cam handle to the locked position.
   d. Interrupt the boot process by pressing `Ctrl-C` when you see Press Ctrl-C for Boot Menu.
   e. Select the option to boot to LOADER.

Restore and verify the system configuration - FAS9500

After completing the hardware replacement, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary..

Step 1: Set and verify the system time after replacing the controller module

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the replacement controller module, verify that all components display the same `HA` state: `ha-config show`

If your system is in…	The HA state for all components should be…
An HA pair	ha
A MetroCluster FC configuration with four or more nodes	mcc
A MetroCluster IP configuration	mccip

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller` *`ha-state`*
3. If the displayed system state of the chassis does not match your system configuration, set the `HA` state for the chassis: `ha-config modify chassis` *`ha-state`*

Recable the system - FAS9500

Continue the replacement procedure by recabling the storage and network conigurations.

Step 1: Recable the system

You must recable the controller module's storage and network connections.

Steps

1. Recable the system.
2. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

> The system ID and disk assignment information reside in the NVRAM module, which is in a module separate from the controller module and not impacted by the controller module replacement.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to

the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* node and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* node is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* node, boot the node, entering `y` if you are prompted to override the system ID due to a system ID mismatch.`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* node console and then, from the healthy node, verify that the new partner system ID has been automatically assigned: storage failover show

   In the command output, you should see a message that the system ID has changed on the impaired node, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> storage failover show
                                    Takeover
Node                Partner         Possible     State Description
------------      ------------      --------
-------------------------------------
node1               node2             false        System ID changed on
partner (Old:
                                                    151759755, New:
151759706), In takeover
node2               node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy node, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`
   c. Wait for the savecore command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *`local-node-name`* `partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys

6. Give back the node:

   a. From the healthy node, give back the replaced node's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* node takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: storage failover show

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* node should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the node: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each node will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the node is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a node on the disaster site.

   This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* node is the current owner of the disks on the disaster site.

  For more information, see Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration topic.

10. If your system is in a MetroCluster configuration, verify that each node is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each node: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy node: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - FAS9500

To complete the replacement procedure and restore your system to full operation, you must recable the storage, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Install licenses for the replacement node in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The licenses keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.
- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

   The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return

and Replacements page for further information.

Replace a DIMM - FAS9500

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

Before you begin

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired node

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

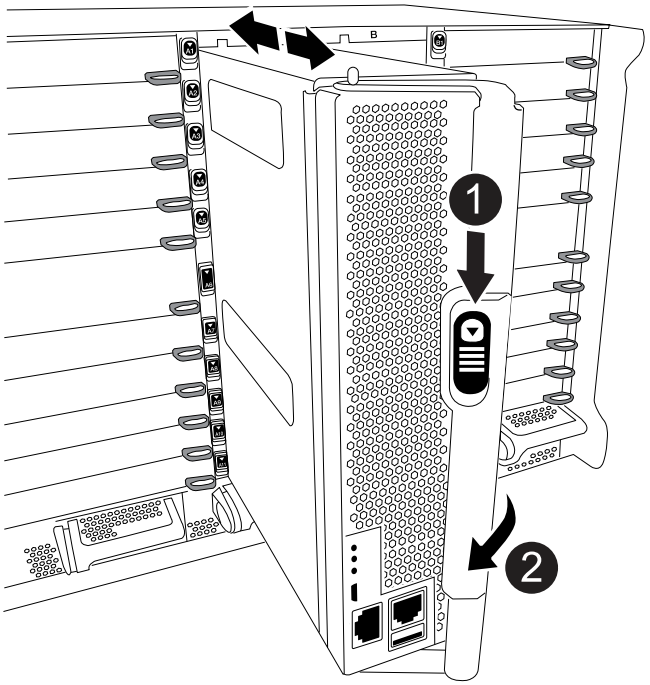

1	Cam handle release button

2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

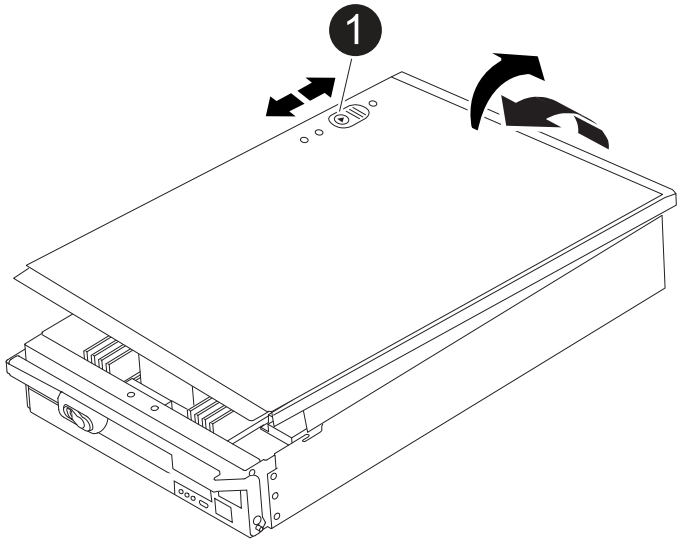

1	Controller module cover locking button

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

> The VER2 controller has fewer DIMM sockets. There is no reduction in the number of DIMMs supported or change in the DIMM socket numbering. When moving the DIMMs to the new controller module, install the DIMMs into the same socket number/location as the impaired controller module. See the FRU map diagram on the VER2 controller module for DIMM socket locations.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.

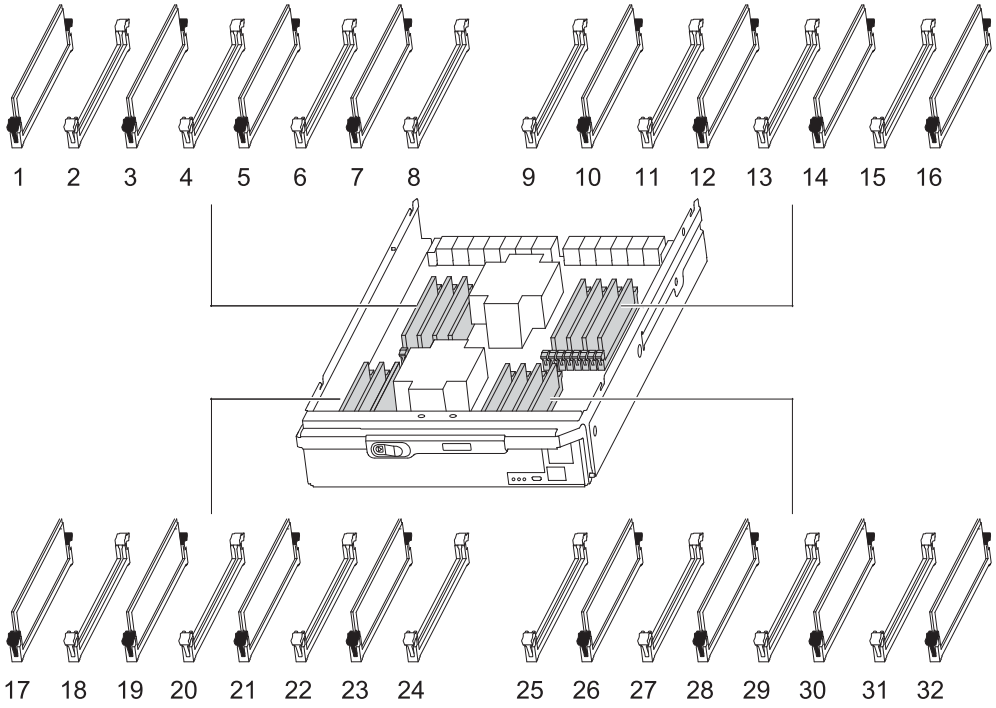

3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

> Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

Animation - Replace DIMMs

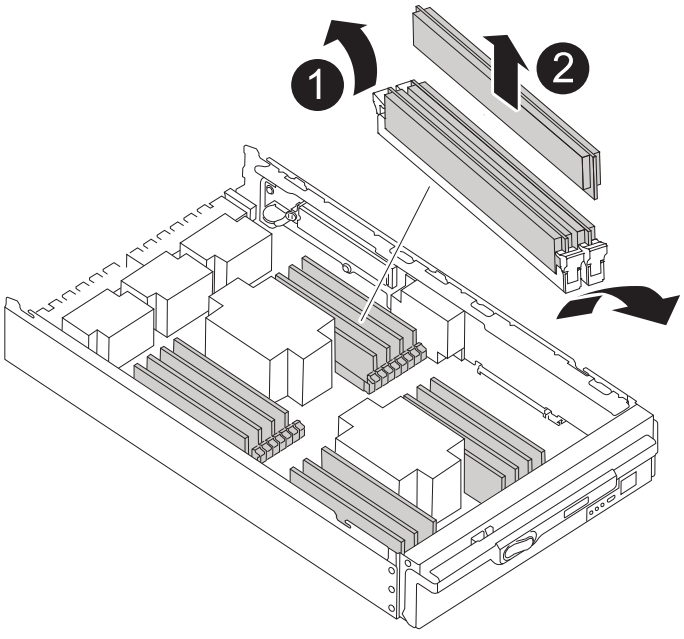

1	DIMM ejector tabs

2	DIMM

4. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

5. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Close the controller module cover.

Step 4: Install the controller

After you install the components into the controller module, you must install the controller module back into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.

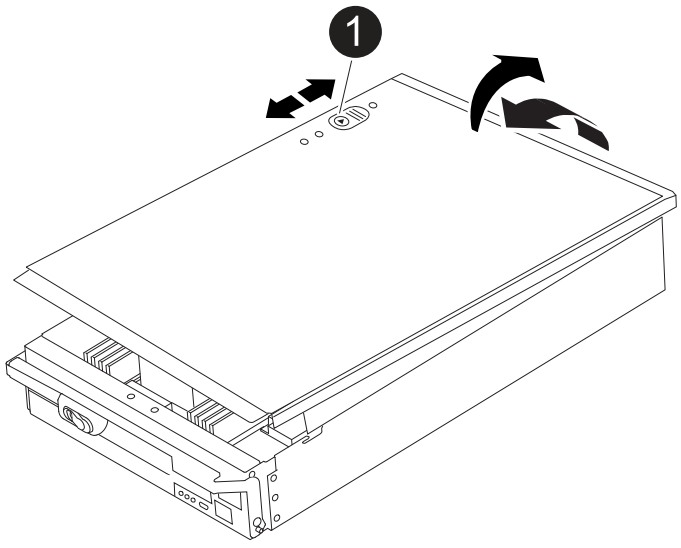

1	Controller module cover locking button

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller

module halfway into the system.

Animation - Install controller

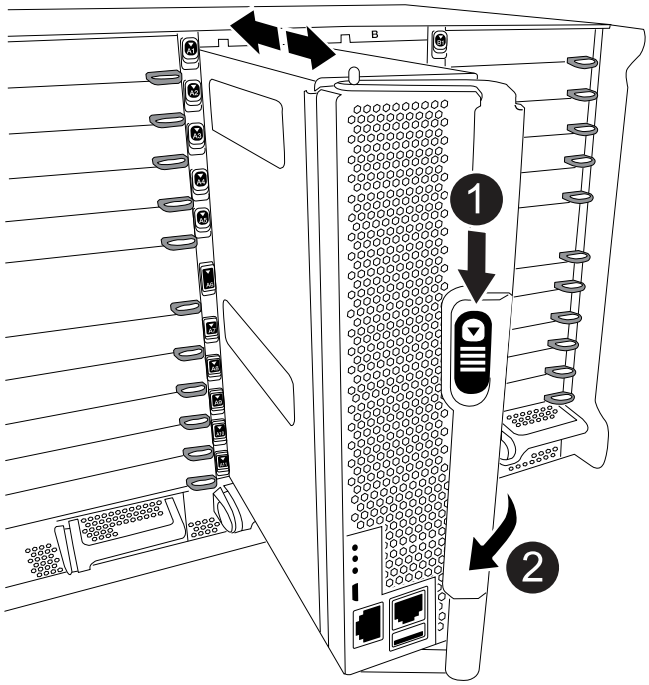

1	Cam handle release button
2	Cam handle

Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the Destage Control Power Module containing the NVRAM11 battery - FAS9500

To hot-swap a destage controller power module (DCPM), which contains the NVRAM11 battery, you must locate the failed DCPM module, remove it from the chassis, and install the replacement DCPM module.

You must have a replacement DCPM module in-hand before removing the failed module from the chassis and it must be replaced within five minutes of removal. Once the DCPM module is removed from the chassis, there is no shutdown protection for the controller module that owns the DCPM module, other than failover to the other controller module.

Step 1: Replace the DCPM module

To replace the DCPM module in your system, you must remove the failed DCPM module from the system and then replace it with a new DCPM module.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel on the front of the system and set it aside.
3. Locate the failed DCPM module in the front of the system by looking for the Attention LED on the module.

   The LED will be steady amber if the module is faulty.

   > The DCPM module must be replaced in the chassis within five minutes of removal or the associated controller will shut down.

4. Press the terra cotta locking button on the module handle, and then slide the DCPM module out of the chassis.

   Animation - Remove/install DCPM

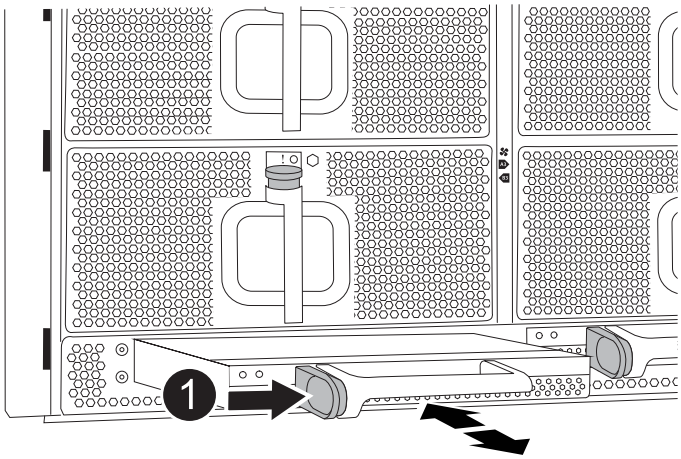

1	DCPM module terra cotta locking button

5. Align the end of the DCPM module with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

   The Amber LED flashes four times upon insertion and the green LED also flashes if the battery is providing a voltage. If it does not flash, it will likely need to be replaced.

Step 2: Dispose of batteries

You must dispose of batteries according to the local regulations regarding battery recycling or disposal. If you cannot properly dispose of batteries, you must return the batteries to NetApp, as described in the RMA instructions that are shipped with the kit.

Safety Information and Regulatory Notices

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a fan - FAS9500

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press the terra cotta button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

   Animation - Remove/install fan

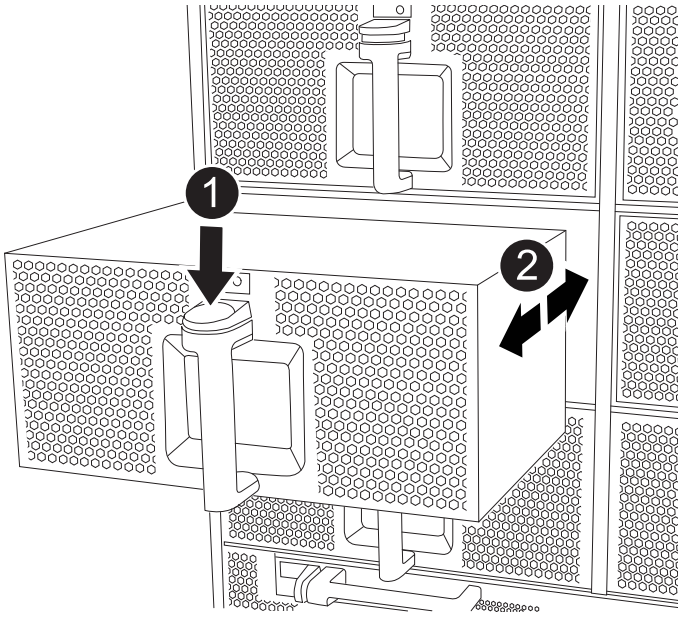

1	Terra cotta release button
2	Slide fan in/out of chassis

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Add an I/O module - FAS9500

You can add an I/O module to your system by either replacing a NIC or storage adapter with a new one in a fully-populated system, or by adding a new NIC or storage adapter into an empty chassis slot in your system.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- To non-disruptively add an I/O module, you must take over the target controller, remove the slot blanking cover in the target slot or remove an existing I/O module, add the new or replacement I/O module, and then

giveback the target controller.

- Make sure that all other components are functioning properly.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster

Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Add the new I/O modules

If the storage system has empty slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Add I/O module to an empty slot

You can add a new I/O module into a storage system with available empty slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the target slot blanking cover:
   a. Depress the lettered and numbered cam latch.
   b. Rotate the cam latch down until it is the open position.
   c. Remove the blanking cover.
3. Install the I/O module:
   a. Align the I/O module with the edges of the slot.
   b. Slide the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
4. If the replacement I/O module is a NIC, cable the module to the data switches.

   Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

5. Reboot the controller from the LOADER prompt: *bye*

   This reinitializes the PCIe cards and other components and reboots the node.

6. Give back the node from the partner node. `storage failover giveback -ofnode target_node_name`
7. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto -giveback true`
8. If you are using slots 3 and/or 7 for networking, use the `storage port modify -node <node name> -port <port name> -mode network` command to convert the slot for networking use.
9. Repeat these steps for controller B.
10. If you installed a storage I/O module, install and cable your SAS shelves, as described in Hot-adding a SAS shelf.

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam latch.

      The cam latch moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

      Animation - Replace an I/O module

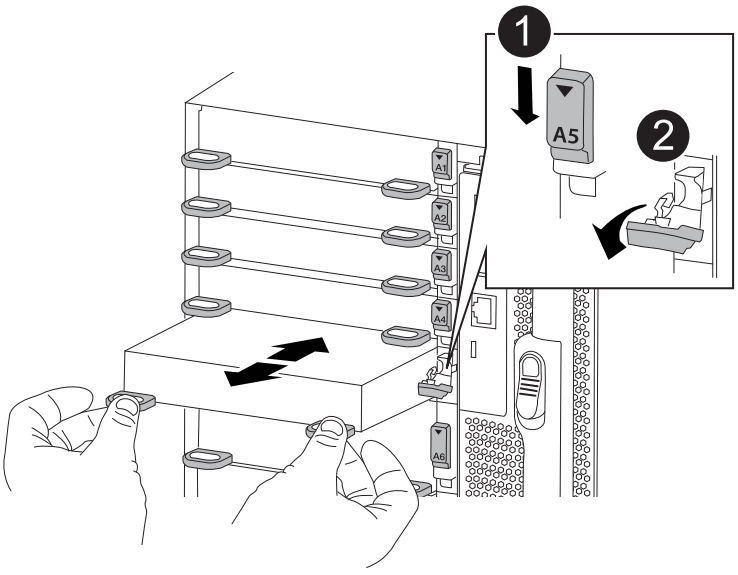

1	Lettered and numbered I/O cam latch

2	I/O cam latch completely unlocked

4. Install the I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Slide the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
5. Repeat the remove and install steps to replace additional modules for controller A.
6. If the replacement I/O module is a NIC, cable the module or modules to the data switches.

   This reinitializes the PCIe cards and other components and reboots the node.

7. Reboot the controller from the LOADER prompt:
   a. Check the version of BMC on the controller: `system service-processor show`
   b. Update the BMC firmware if needed: `system service-processor image update`
   c. Reboot the node: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.

   If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement

8. Give back the node from the partner node. `storage failover giveback -ofnode target_node_name`
9. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto -giveback true`
10. If you added:

If I/O module is a…	Then…
NIC module in slots 3 or 7	Use the `storage port modify -node *<node name> -port *<port name> -mode network` command for each port.
Storage module	Install and cable your SAS shelves, as described in Hot-adding a SAS shelf.

11. Repeat these steps for controller B.

Replace an I/O module - FAS9500

To replace an I/O module, you must perform a specific sequence of tasks.

- You can use this procedure with all versions of ONTAP supported by your system.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired node

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace I/O modules

To replace an I/O module, locate it within the chassis and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

3. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

      Animation - Remove/install I/O module

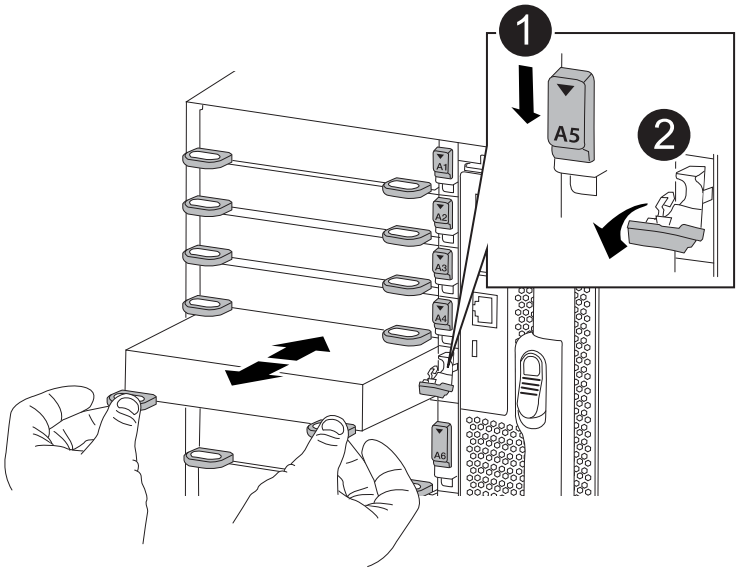

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Install the replacement I/O module into the chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.
6. Recable the I/O module, as needed.

Step 3: Reboot the controller after I/O module replacement

After you replace an I/O module, you must reboot the controller module.

If the new I/O module is not the same model as the failed module, you must first reboot the BMC.

Steps

1. Reboot the BMC if the replacement module is not the same model as the old module:
   a. From the LOADER prompt, change to advanced privilege mode: `priv set advanced`
   b. Reboot the BMC: `sp reboot`
2. From the LOADER prompt, reboot the node: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.

3. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs, convert these ports to 10 GbE connections by using the `nicadmin convert` command from Maintenance mode. See Convert 40GbE NIC ports into multiple 10GbE ports for 10GbE connectivity for more information.

   Be sure to exit Maintenance mode after completing the conversion.

4. Return the node to normal operation: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an LED USB module - FAS9500

The LED USB module provides connectivity to console ports and system status. Replacement of this module does not require tools and does not interrupt service.

Step 1: Replace the LED USB module

Steps

1. Remove the old LED USB module:

   Animation - Remove/install LED-USB module

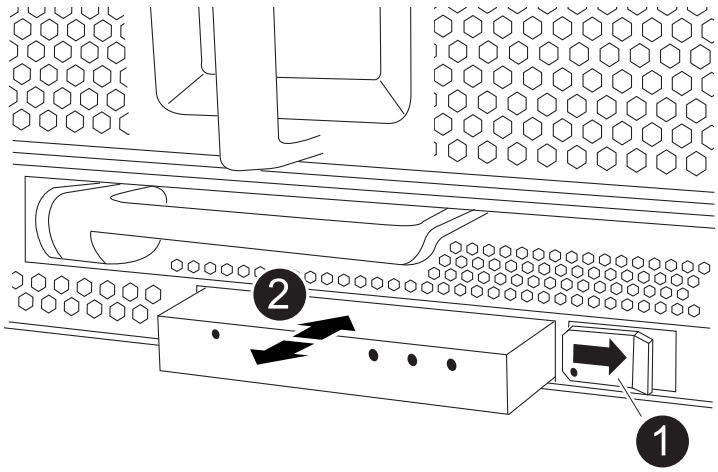

1	Locking button
2	USB LED module

a. With the bezel removed, locate the LED USB module at the front of the chassis, on the bottom left side.
b. Slide the latch to partially eject the module.
c. Pull the module out of the bay to disconnect it from the midplane. Do not leave the slot empty.

2. Install the new LED USB module:
   a. Align the module to the bay with the notch in the corner of the module positioned near the slider latch on the chassis. The bay will prevent you from installing the module upside down.
   b. Push the module into the bay until it is fully seated flush with the chassis.

      There is an audible click when the module is secure and connected to the midplane.

Step 2: Return the failed component

1. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVRAM module and/or NVRAM DIMMs - FAS9500

The NVRAM module consists of the NVRAM11 and DIMMs. You can replace a failed NVRAM module or the DIMMs inside the NVRAM module. To replace a failed NVRAM module, you must remove it from the chassis, move the DIMMs to the replacement module, and install the replacement NVRAM module into the chassis.

To replace and NVRAM DIMM, you must remove the NVRAM module from the chassis, replace the failed DIMM in the module, and then reinstall the NVRAM module.

About this task

Because the system ID is derived from the NVRAM module, if replacing the module, disks belonging to the system are reassigned to a new system ID.

Before you begin

- All disk shelves must be working properly.
- If your system is in an HA pair, the partner controller must be able to take over the controller associated with the NVRAM module that is being replaced.
- This procedure uses the following terminology:
  - The impaired controller is the controller on which you are performing maintenance.
  - The healthy controller is the HA partner of the impaired controller.
- This procedure includes steps for automatically reassigning disks to the controller module associated with the new NVRAM module. You must reassign the disks when directed to in the procedure. Completing the disk reassignment before giveback can cause issues.
- You must replace the failed component with a replacement FRU component you received from your provider.

- You cannot change any disks or disk shelves as part of this procedure.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace the NVRAM module

To replace the NVRAM module, located it in slot 6 in the chassis and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the target NVRAM module from the chassis:
   a. Depress the lettered and numbered cam latch.

      The cam latch moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

The NVRAM module disengages from the chassis and moves out a few inches.

c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

Animation - Replace the NVRAM module

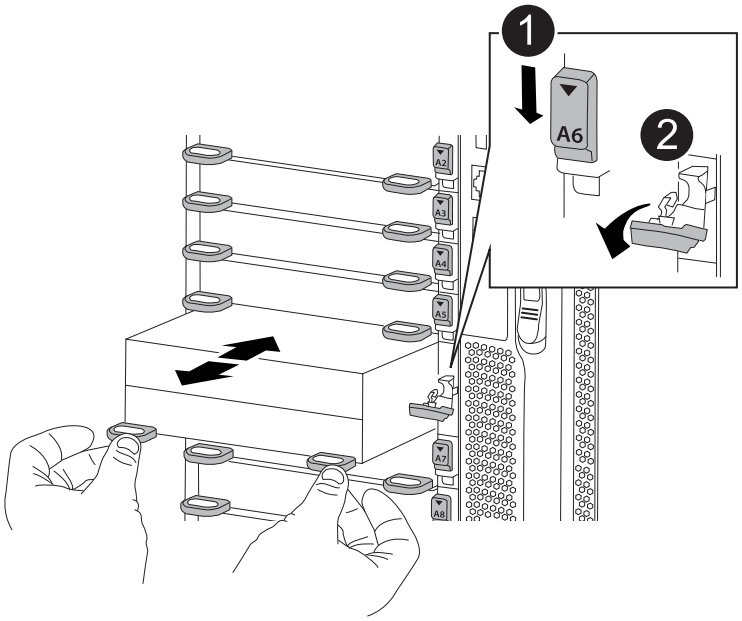

1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

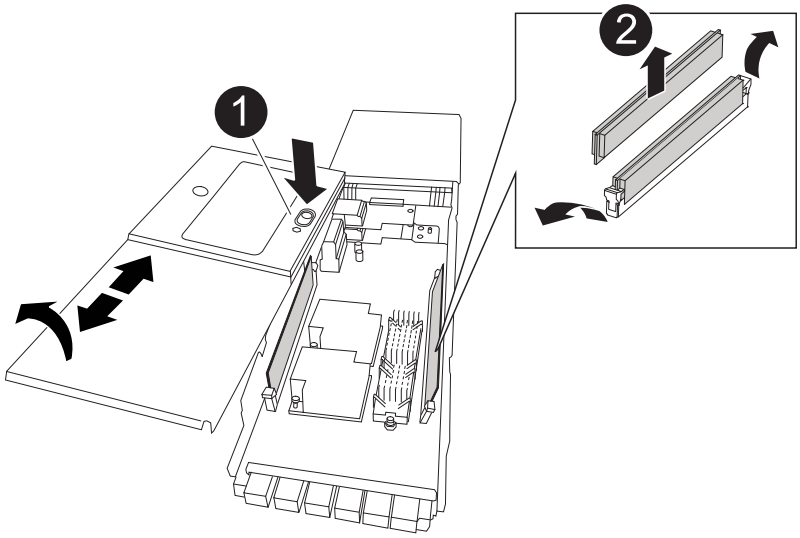

1	Cover locking button
2	DIMM and DIMM ejector tabs

4. Remove the DIMMs, one at a time, from the old NVRAM module and install them in the replacement NVRAM module.
5. Close the cover on the module.
6. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

Step 3: Replace a NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, you must remove the NVRAM module, open the module, and then replace the target DIMM.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the target NVRAM module from the chassis:
   a. Depress the lettered and numbered cam latch.

      The cam latch moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The NVRAM module disengages from the chassis and moves out a few inches.

   c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

      Animation - Replace the NVRAM module

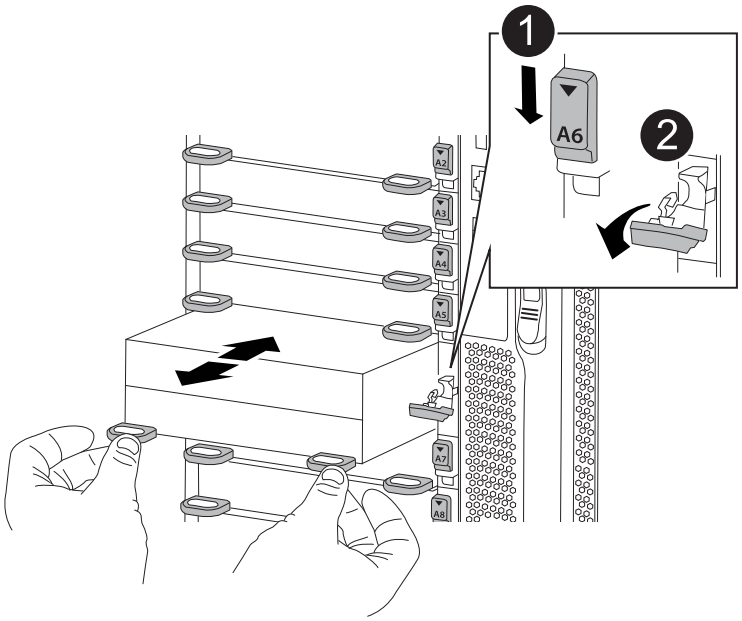

1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

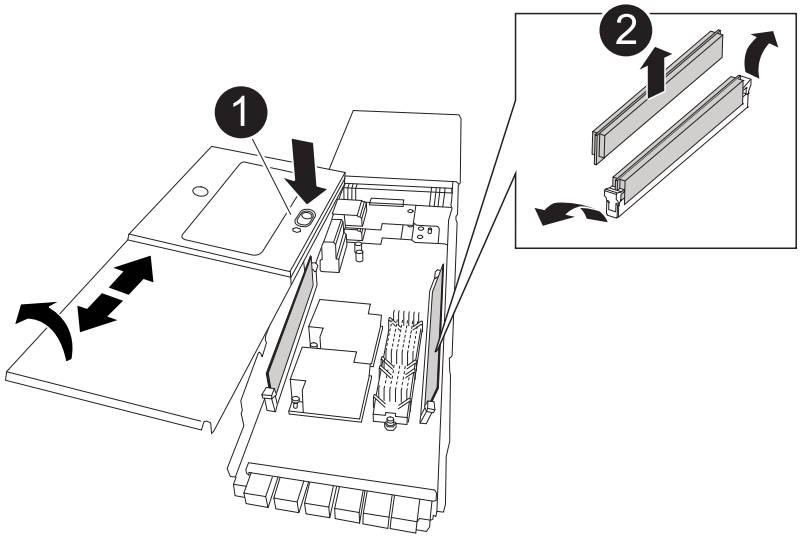

1	Cover locking button
2	DIMM and DIMM ejector tabs

4. Locate the DIMM to be replaced inside the NVRAM module, and then remove it by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.

5. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
6. Close the cover on the module.
7. Install the NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

Step 4: Reboot the controller after FRU replacement

After you replace the FRU, you must reboot the controller module.

To boot ONTAP from the LOADER prompt, enter `bye`.

Step 5: Reassigning disks

You must confirm the system ID change when you boot the replacement controller and then verify that the change was implemented.

Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the replacement controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the replacement controller, boot the controller and entering y if you are prompted to override the system ID due to a system ID mismatch.
3. Wait until the Waiting for giveback… message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1:> storage failover show
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. Give back the controller:

   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The replacement controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

5. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the replacement controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 151759706:

```
node1:> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        151759706  151759706  -
151759706 Pool0
1.0.1  aggr0_1  node1 node1           151759706  151759706  -
151759706 Pool0
.
.
.
```

6. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

7. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The replacement controller is the current owner of the disks on the disaster site.

  See Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration for more information.

8. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

9. Verify that the expected volumes are present for each controller: `vol show -node node-name`
10. If storage encryption is enabled, you must restore functionality.
11. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - FAS9500

Swapping out a power supply involves turning off, disconnecting, and removing the power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

About this task

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- There are four power supplies in the system.

- Power supplies are auto-ranging.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
4. Press and hold the terra cotta button on the power supply handle, and then pull the power supply out of the chassis.

> When removing a power supply, always use two hands to support its weight.

Animation - Remove/install PSU

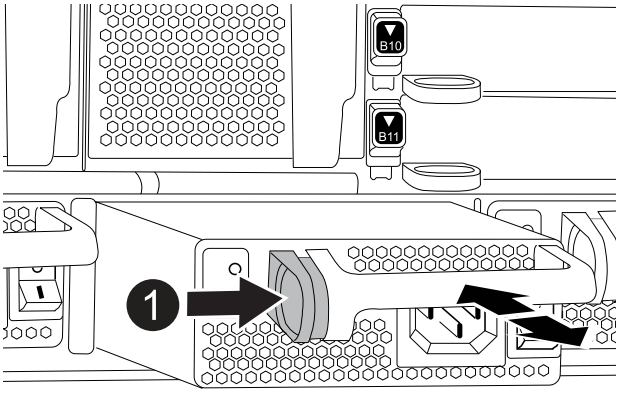

1	Locking button

5. Make sure that the on/off switch of the new power supply is in the Off position.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

> Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply.
   b. Secure the power cable to the power supply using the power cable retainer.

      Once power is restored to the power supply, the status LED should be green.

8. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

   The green power LED lights when the PSU is fully inserted into the chassis and the amber attention LED flashes initially, but turns off after a few moments.

9. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - FAS9500

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired node

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove controller module

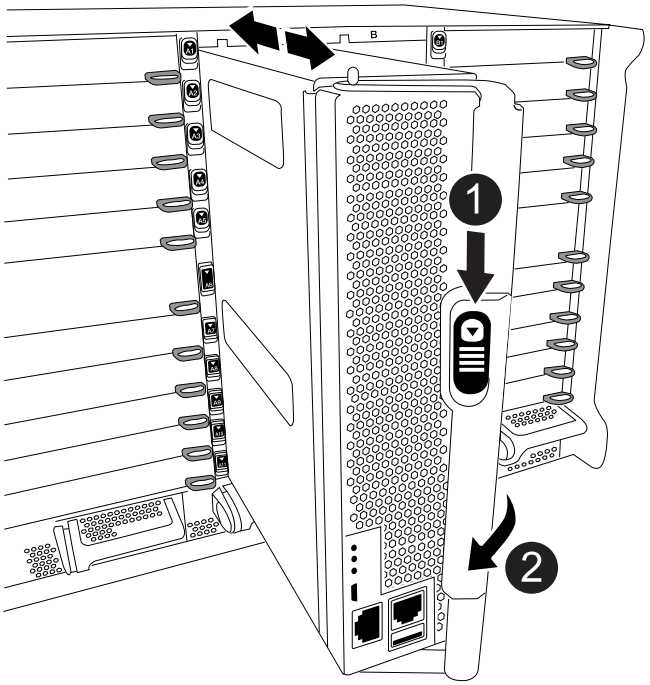

1	Cam handle release button

2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

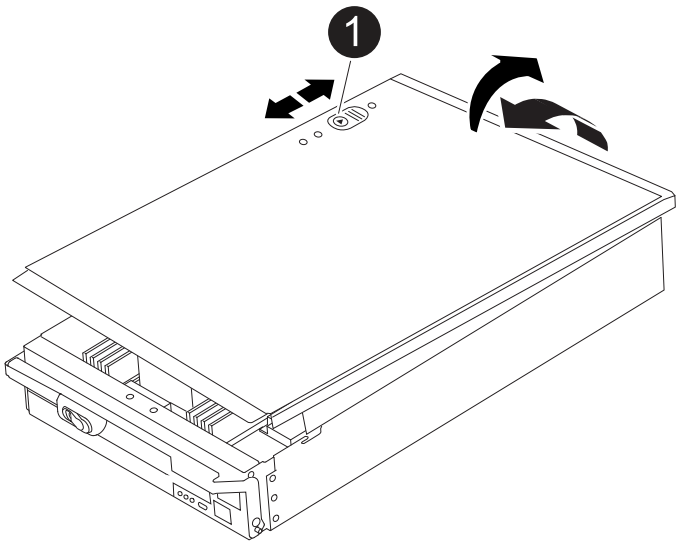

1	Controller module cover locking button

Step 3: Replace the RTC battery

To replace the RTC battery, you must locate the failed battery in the controller module, remove it from the holder, and then install the replacement battery in the holder.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

   Animation - Replace RTC battery

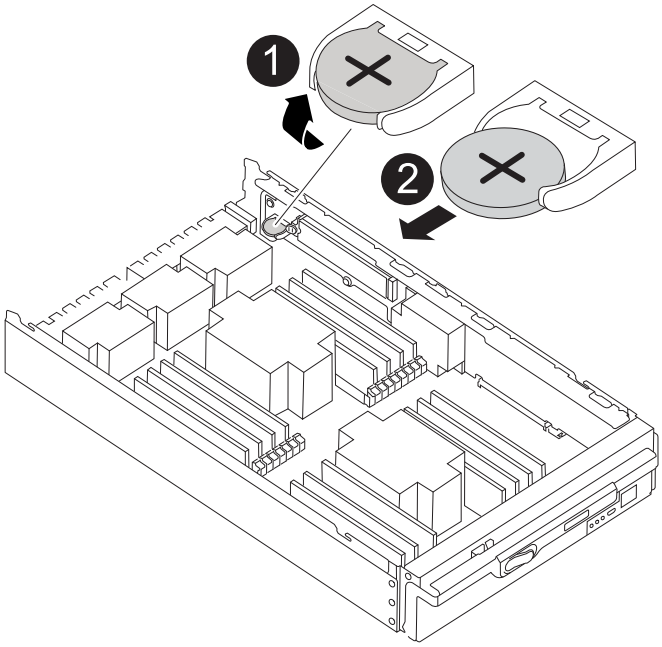

1	Rotate battery up
2	Slide battery out from housing

Steps

1. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

2. Remove the replacement battery from the antistatic shipping bag.
3. Locate the empty battery holder in the controller module.
4. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
5. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
6. Reinstall the controller module cover.

Step 4: Reinstall the controller module and set time/date

After you replace the RTC battery, you must reinstall the controller module. If the RTC battery has been left out of the controller module for more than 10 minutes, you may have to reset the time and date.

Steps

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.

If your system stops at the boot menu, select the option for "Reboot node" and respond y when prompted,
then boot to LOADER by pressing `Ctrl-C`.

1. Reset the time and date on the controller:
   a. Check the date and time on the healthy node with the `show date` command.
   b. At the LOADER prompt on the target node, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target node.
2. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the node reboot.
3. Return the node to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

End-of-availability systems

AFF A200 systems

Install and setup

Cluster configuration worksheet - AFF A200

You can use the Cluster Configuration Worksheet to gather and record your site-specific IP addresses and other information required when configuring an ONTAP cluster.

Start here: Choose your installation and setup experience

You can choose from different content formats to guide you through installing and setting up your new storage system.

- AFF A200 Installation and Setup Instructions

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

Installation and setup PDF poster - AFF A200

You can use the AFF A200 Installation and Setup Instructions poster to install and set up your new system. The PDF poster provides step-by-step instructions with live links to additional content.

Maintain

Maintain AFF A200 hardware

For the AFF A200 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Caching module

You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

NVEM battery

A battery is included with a controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF A200

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

What you'll need

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

Before you begin

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the var file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the var file system.
  - For disruptive replacement, you do not need a network connection to restore the var file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
  - The *impaired* controller is the controller on which you are performing maintenance.
  - The *healthy* controller is the HA partner of the impaired controller.

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the impaired controller - AFF A200

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller displays…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Replace the boot media - AFF A200

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

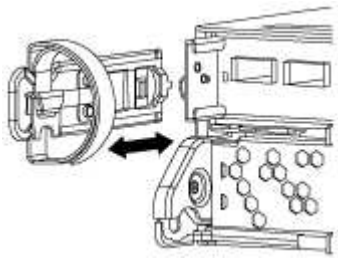

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

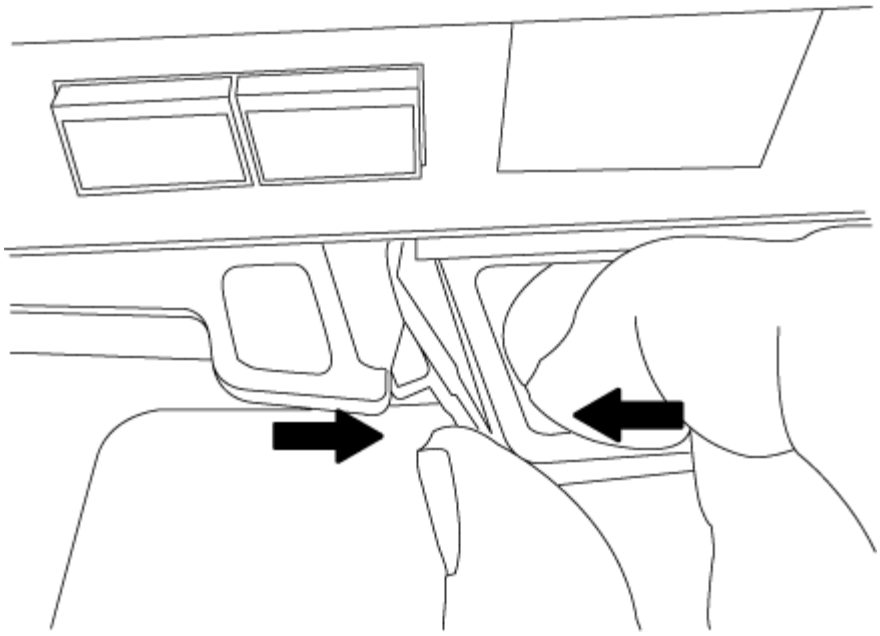

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

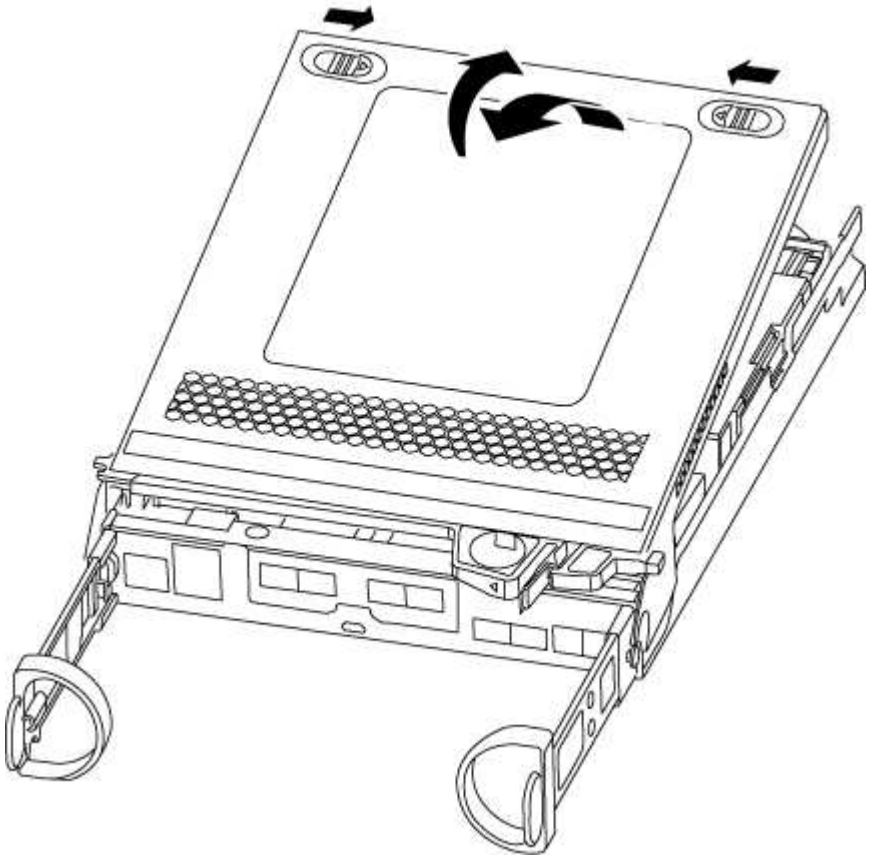

Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the boot media using the following illustration or the FRU map on the controller module:
3. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

4. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
5. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

6. Push the boot media down to engage the locking button on the boot media housing.
7. Close the controller module cover.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

What you'll need

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the

closed position, and then tighten the thumbscrew.

The controller begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

6. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

   The system begins to boot and stops at the LOADER prompt.

7. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

     The target port you configure is the target port you use to communicate with the impaired controller from the healthy controller during var file system restore with a network connection. You can also use the e0M port in this command.

   - If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
     - `filer_addr` is the IP address of the storage system.
     - `netmask` is the network mask of the management network that is connected to the HA partner.
     - `gateway` is the gateway for the network.
     - `dns_addr` is the IP address of a name server on your network.
     - `dns_domain` is the Domain Name System (DNS) domain name.

       If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

   Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

Boot the recovery image - AFF A200

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the `var` file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `saveenv` command.
5. The next depends on your system configuration:
   - If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed
   - If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.
6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.

8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore encryption - AFF A200

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
**************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
**************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
**********************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
**********************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF A200

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF A200

To replace the chassis, move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

What you'll need

All other components in the system must be functioning properly; if not, contact technical support.

About this task

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving all drives and controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A200

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - AFF A200

Move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Move the power supply

Move the power supply from the old chassis to the replacement chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.
4. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

> Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

Step 2: Remove the controller module

Remove the controller module or modules from the old chassis.

Steps

1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

2. Remove and set aside the cable management devices from the left and right sides of the controller module.

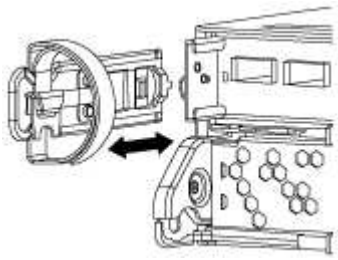

3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

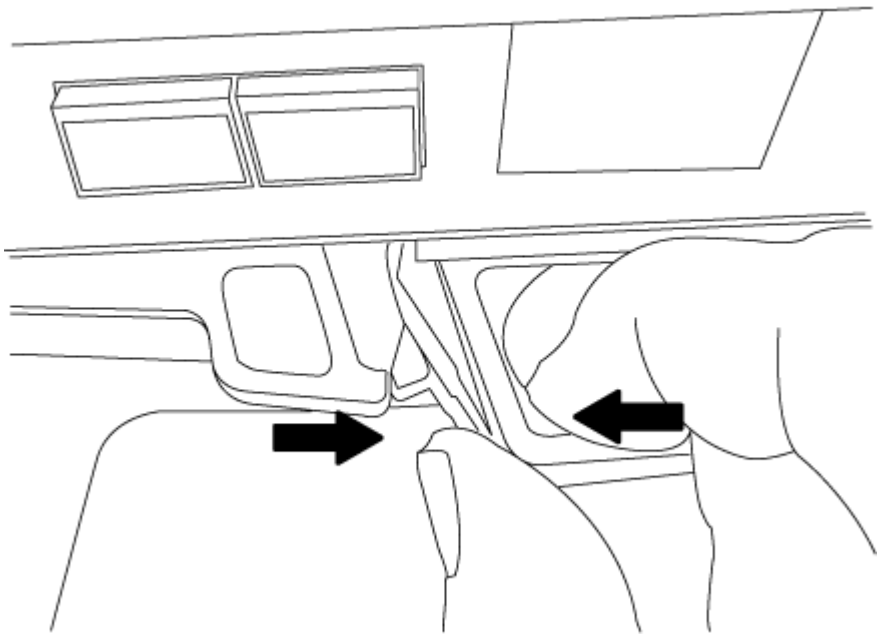

4. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 3: Move drives to the new chassis

Move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

Steps

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.
5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.
6. Repeat the process for the remaining drives in the system.

Step 4: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

Steps

1. Remove the screws from the chassis mount points.
2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller

After you install the controller module and any other components into the new chassis, you must boot your system.

About this task

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Repeat the preceding steps if there is a second controller to install in the new chassis.
4. Complete the installation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Repeat the preceding steps for the second controller module in the new chassis.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reinstall the blanking panel and then go to the next step.

5. Connect the power supplies to different power sources, and then turn them on.
6. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

      If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

   b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - AFF A200

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller module

Overview of controller module replacement - AFF A200

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

What you'll need

- All drive shelves must be working properly.

- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired node”).

About this task

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must replace a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* node so that the *replacement* node will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* node is the controller that is being replaced.
  - The *replacement* node is the new controller that is replacing the impaired controller.
  - The *healthy* node is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - AFF A200

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *impaired_node_name* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Replace the controller module hardware - AFF A200

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove controller module

To replace the controller module, you must first remove the old controller module from the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

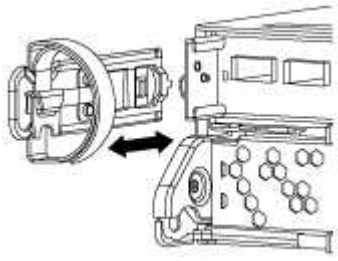

4. If you left the SFP modules in the system after removing the cables, move them to the new controller module.
5. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

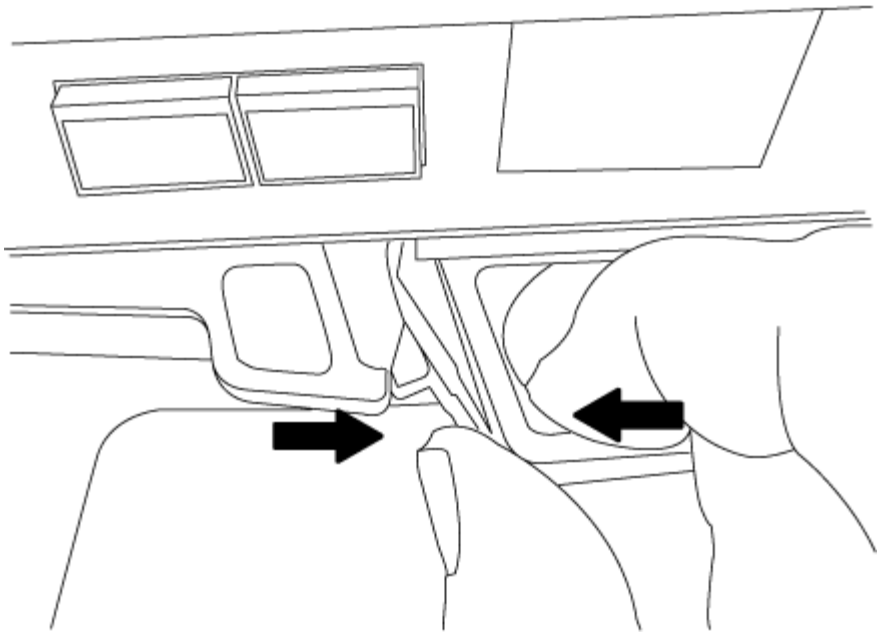

6. Turn the controller module over and place it on a flat, stable surface.
7. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

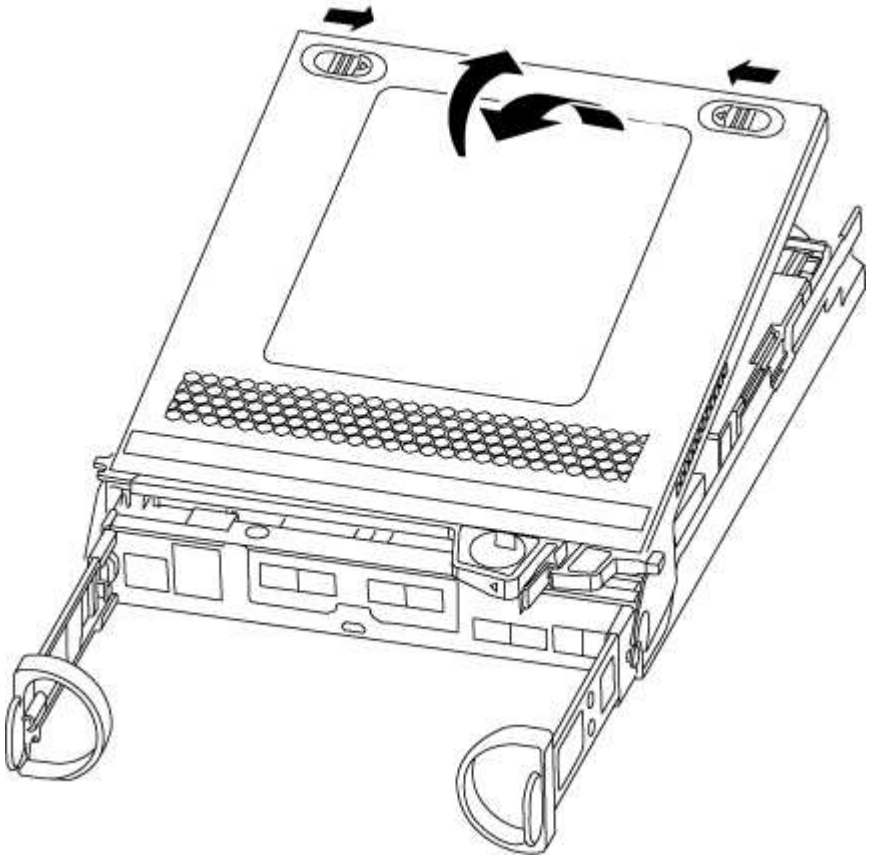

Step 2: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller module and insert it in the new controller module.

Steps

1. Locate the boot media using the following illustration or the FRU map on the controller module:
2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the NVMEM battery

To move the NVMEM battery from the old controller module to the new controller module, you must perform a specific sequence of steps.

Steps

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

     

     The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

     - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
     - If the LED is on and power is on, unwritten data is stored on NVMEM.

       This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Locate the NVMEM battery in the controller module.

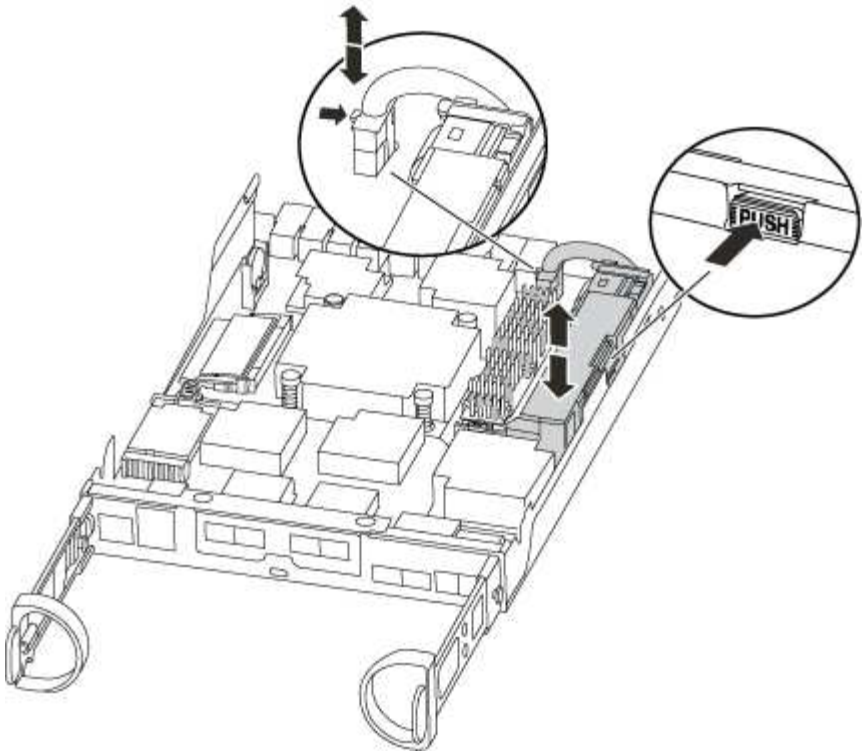

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Loop the battery cable around the cable channel on the side of the battery holder.
7. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
8. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.

Step 4: Move the DIMMs

To move the DIMMs, you must follow the directions to locate and move them from the old controller module into the replacement controller module.

About this task

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

Steps

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

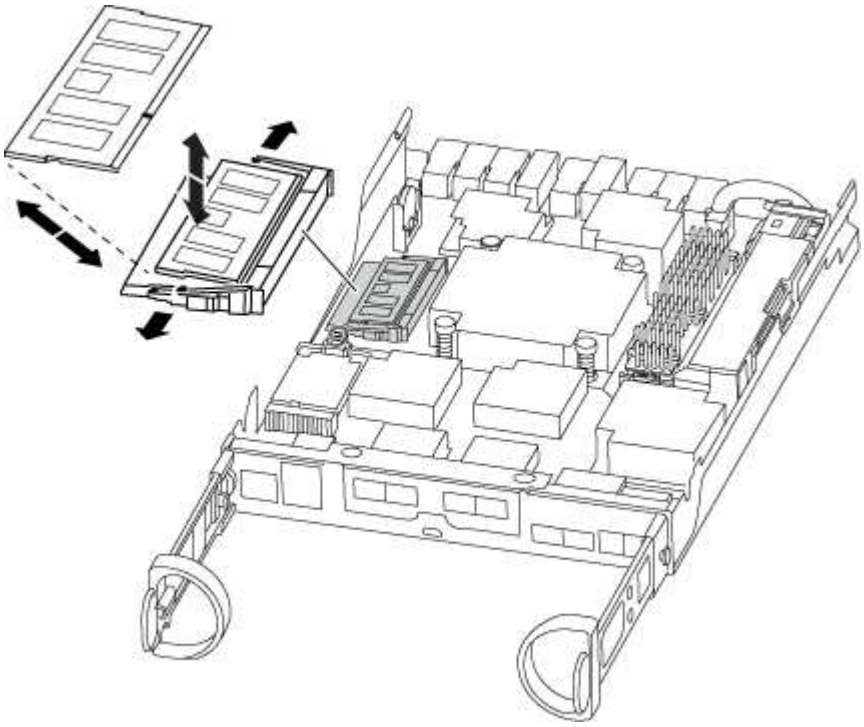

4. Repeat these steps to remove additional DIMMs as needed.
5. Verify that the NVMEM battery is not plugged into the new controller module.
6. Locate the slot where you are installing the DIMM.
7. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Repeat these steps for the remaining DIMMs.
9. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

   Make sure that the plug locks down onto the controller module.

Step 5: Install the controller

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

About this task

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

> The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   > Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   > You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. When you see the message `Press Ctrl-C for Boot Menu`, press `Ctrl-C` to interrupt the boot process. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. Select the option to boot to Maintenance mode from the displayed menu.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process, and then press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. From the boot menu, select the option for Maintenance mode.

During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
  You can safely respond `y` to these prompts.

Restore and verify the system configuration - AFF A200

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The value for HA-state can be one of the following:

   - ha
   - non-ha

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
3. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF A200

Continue the replacement procedure by re-cabling the storage and confirming disk reassignment.

Step 1: Re-cable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.

b. Enter the information for the target system, and then click Collect Data.

c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.

d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks. You must use the correct procedure for your configuration.

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

About this task

This procedure applies only to systems running ONTAP in an HA pair.

Steps

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch.`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

b. Save any coredumps: `system node run -node local-node-name partner savecore`

c. Wait for the `savecore`command to complete before issuing the giveback.

You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

d. Return to the admin privilege level: `set -privilege admin`

5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:

a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

The *replacement* controller takes back its storage and completes booting.

If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. Verify that the expected volumes are present for each controller: `vol show -node node-name`
9. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Option 2: Manually reassign the system ID on a stand-alone system in ONTAP

In a stand-alone system, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems that are in a stand-alone configuration.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by pressing Ctrl-C, and then select the option to boot to Maintenance mode from the displayed menu.
2. You must enter `Y` when prompted to override the system ID due to a system ID mismatch.
3. View the system IDs: `disk show -a`
4. You should make a note of the old system ID, which is displayed as part of the disk owner column.

   The following example shows the old system ID of 118073209:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                    POOL   SERIAL NUMBER  HOME
--------    -------------            -----  -------------  -------------
disk_name    system-1  (118073209)   Pool0  J8XJE9LC       system-1
(118073209)
disk_name    system-1  (118073209)   Pool0  J8Y478RC       system-1
(118073209)
.
.
.
```

5. Reassign disk ownership by using the system ID information obtained from the disk show command: `disk reassign -s old system ID disk reassign -s 118073209`
6. Verify that the disks were assigned correctly: `disk show -a`

   The disks belonging to the replacement node should show the new system ID. The following example now show the disks owned by system-1 the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                    POOL   SERIAL NUMBER  HOME
--------    -------------            -----  -------------  -------------
disk_name    system-1  (118065481)   Pool0  J8Y0TDZC       system-1
(118065481)
disk_name    system-1  (118065481)   Pool0  J8Y0TDZC       system-1
(118065481)
.
.
.
```

7. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
8. Boot the node: `boot_ontap`

Complete system restoration - AFF A200

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the

failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver`

`* -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A200

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

About this task

- All other components in the system must be functioning properly; if not, you must contact technical support.
- You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

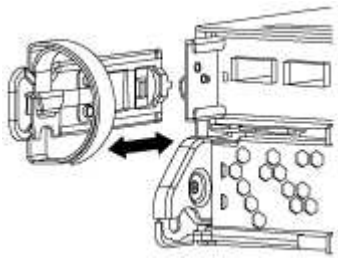

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

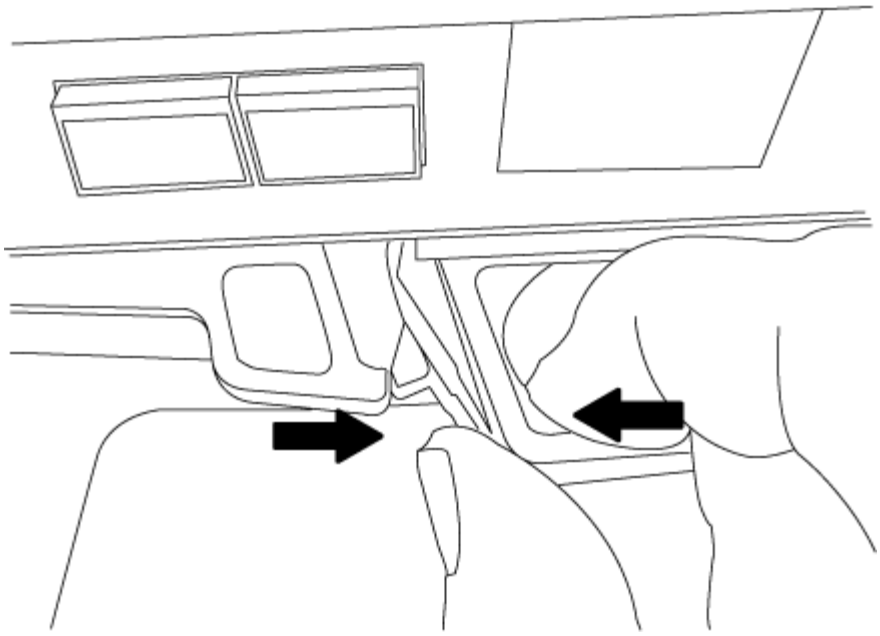

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

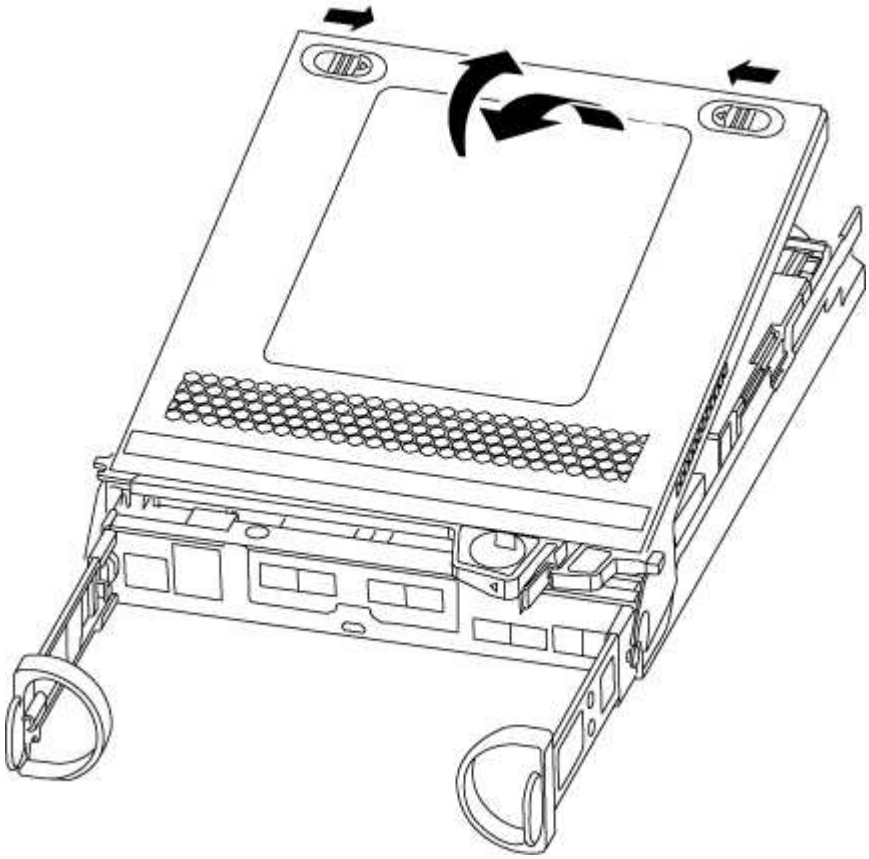

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

About this task

If you are replacing a DIMM, you need to remove it after you have unplugged the NVMEM battery from the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED on the controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

   

3. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
4. If the NVMEM LED is flashing, there is data in the NVMEM and you must disconnect the battery to clear the memory:
   a. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

   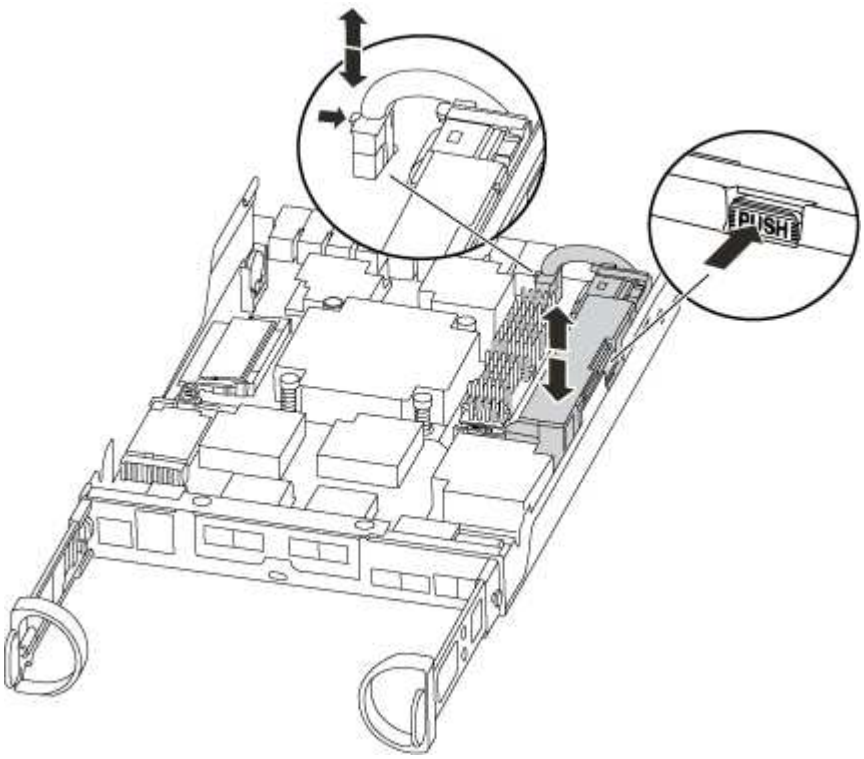

   b. Confirm that the NVMEM LED is no longer lit.
   c. Reconnect the battery connector.
5. Return to step 2 of this procedure to recheck the NVMEM LED.
6. Locate the DIMMs on your controller module.
7. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper

orientation.

8. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   The number and placement of system DIMMs depends on the model of your system.

   The following illustration shows the location of system DIMMs:

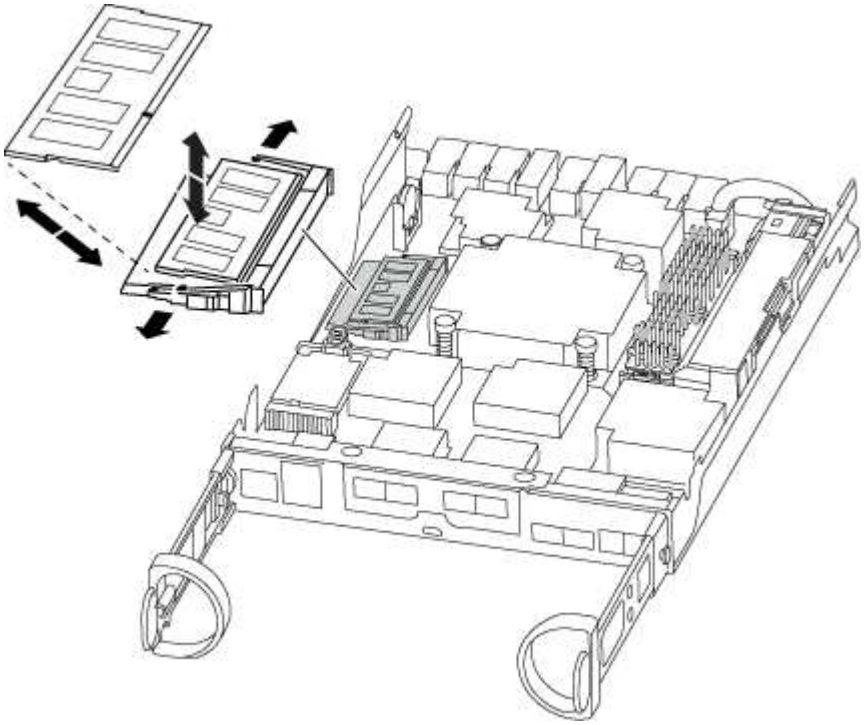

9. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

10. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

    The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

    Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

11. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
12. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

    Make sure that the plug locks down onto the controller module.

13. Close the controller module cover.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, then turn on the power to start the boot process.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - AFF A200

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  > Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace the NVMEM battery - AFF A200

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

About this task

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

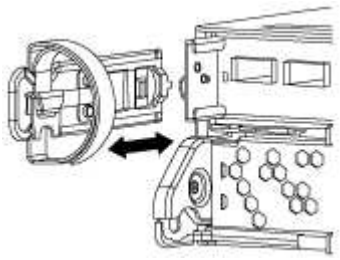

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

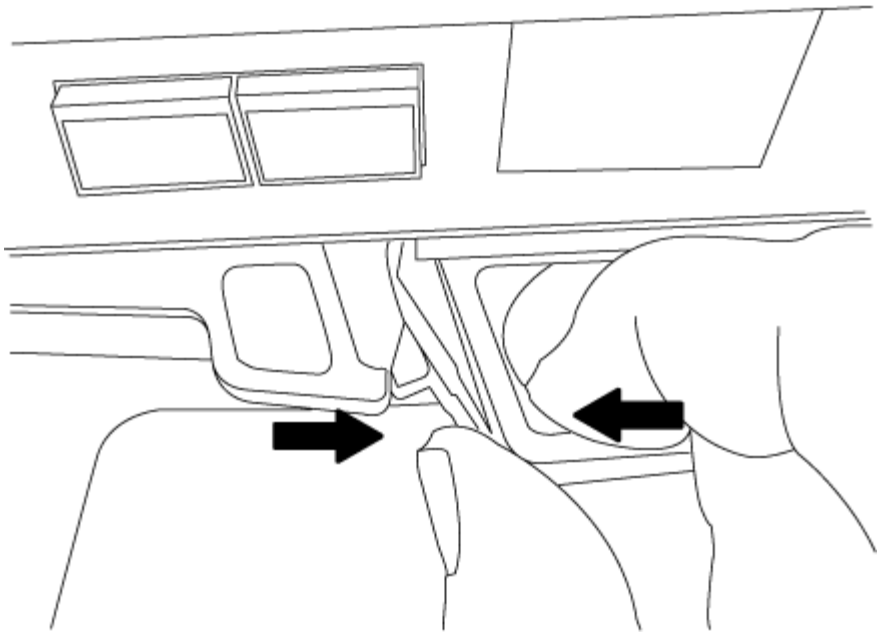

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

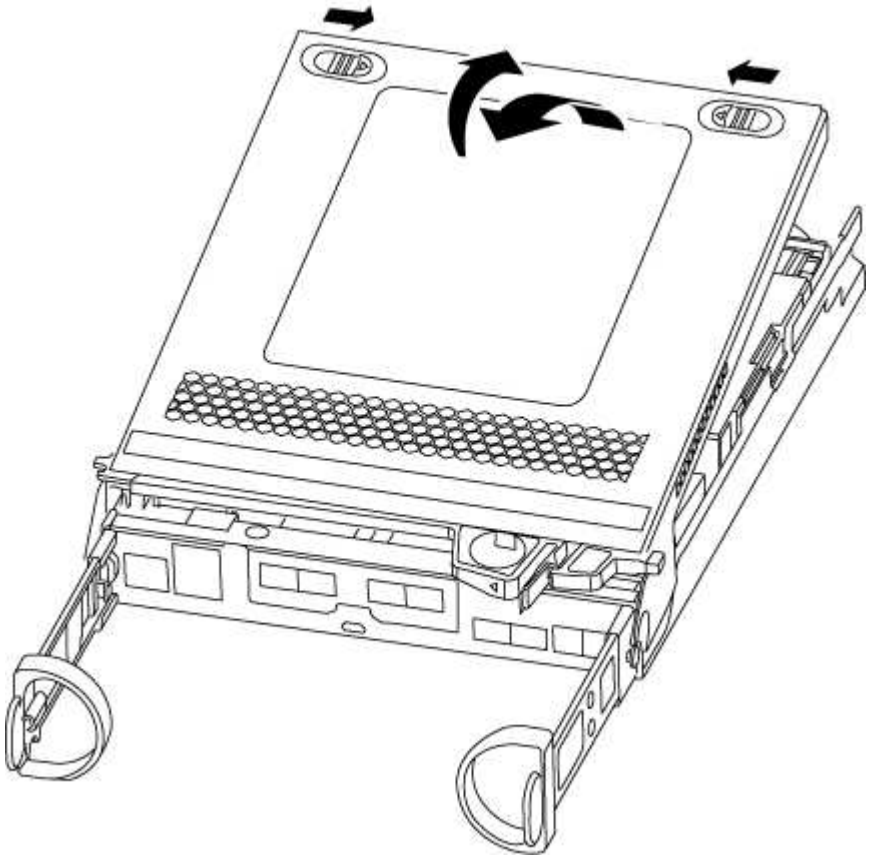

Step 3: Replace the NVMEM battery

To replace the NVMEM battery in your system, you must remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

3. Locate the NVMEM battery in the controller module.

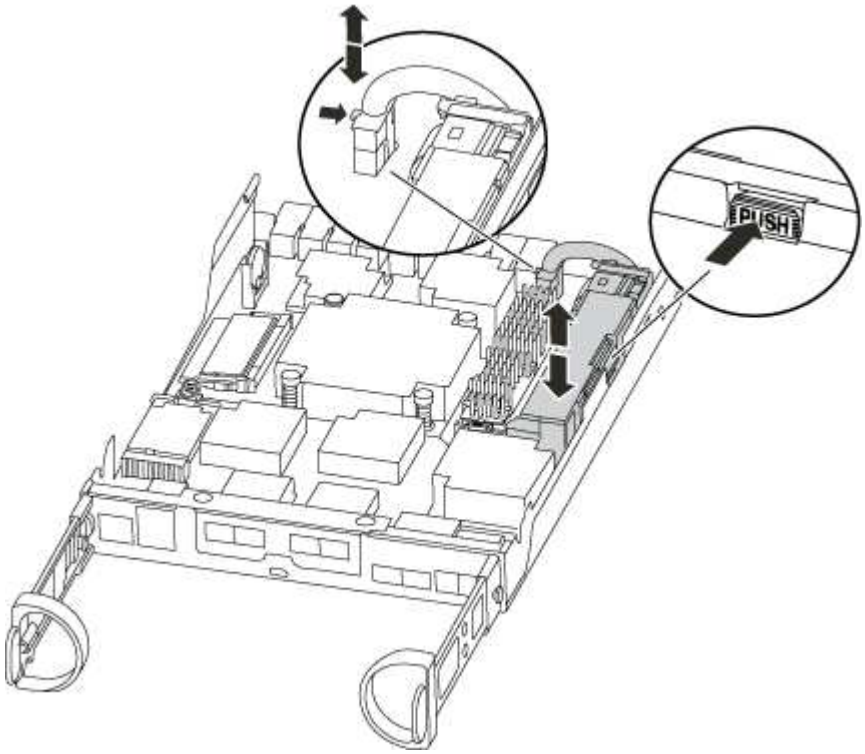

4. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
5. Remove the battery from the controller module and set it aside.
6. Remove the replacement battery from its package.

7. Loop the battery cable around the cable channel on the side of the battery holder.
8. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
9. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
10. Plug the battery plug back into the controller module.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - AFF A200

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

What you'll need

All other components in the system must be functioning properly; if not, you must contact technical support.

About this task

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> Cooling is integrated with the power supply, so you must replace the power supply within two minutes of removal to prevent overheating due to reduced airflow. Because the chassis provides a shared cooling configuration for the two HA nodes, a delay longer than two minutes will shut down all controller modules in the chassis. If both controller modules do shut down, make sure that both power supplies are inserted, turn both off for 30 seconds, and then turn both on.

- The number of power supplies in the system depends on the model.
- Power supplies are auto-ranging.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.

3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.

   If you have an AFF A200 system, a plastic flap within the now empty slot is released to cover the opening and maintain air flow and cooling.

5. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.
7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
9. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.

10. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

    The power supply LEDs are lit when the power supply comes online.

11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A200

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

About this task

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

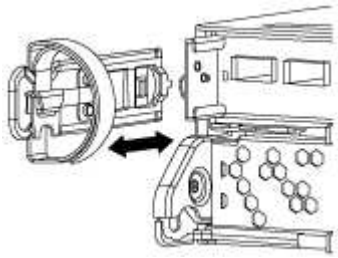

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

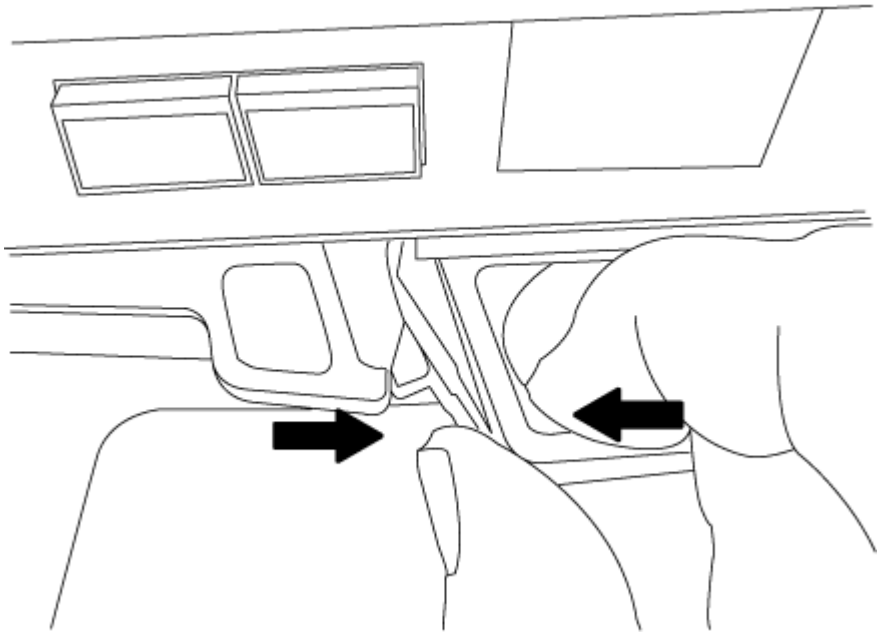

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

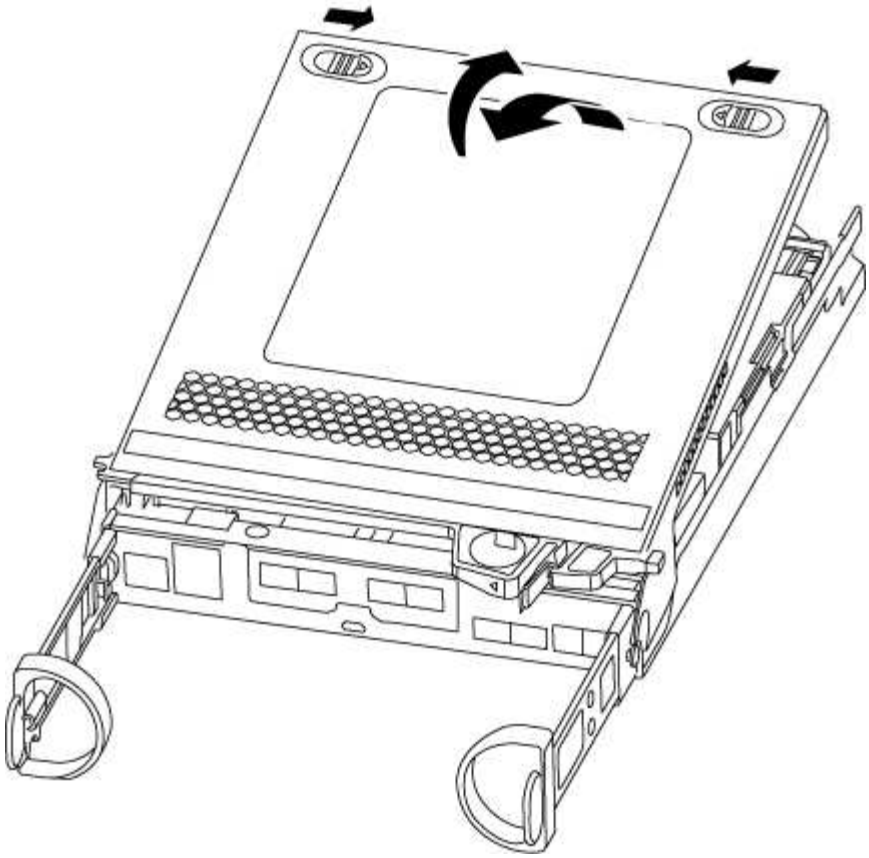

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

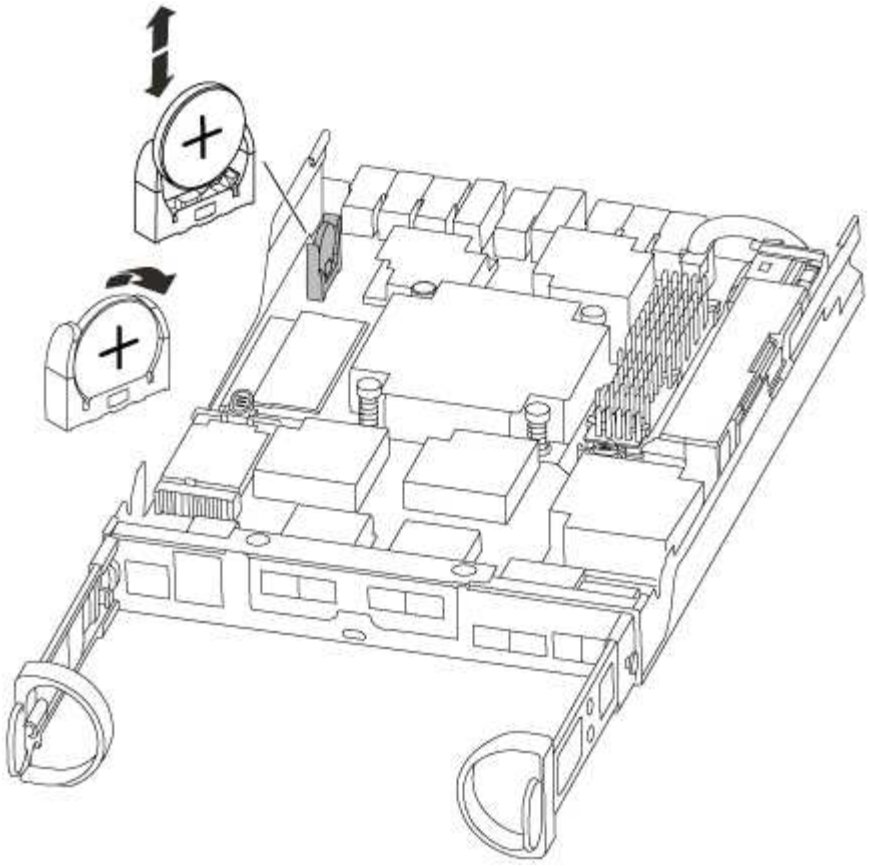

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and set time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

Steps

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF A220 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick guide - AFF A220

This page gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Access the *Installation and Setup Instructions* PDF poster:

AFF A220/FAS2700 Systems Installation and Setup Instructions

Video steps - AFF A220

The following video shows how to install and cable your new system.

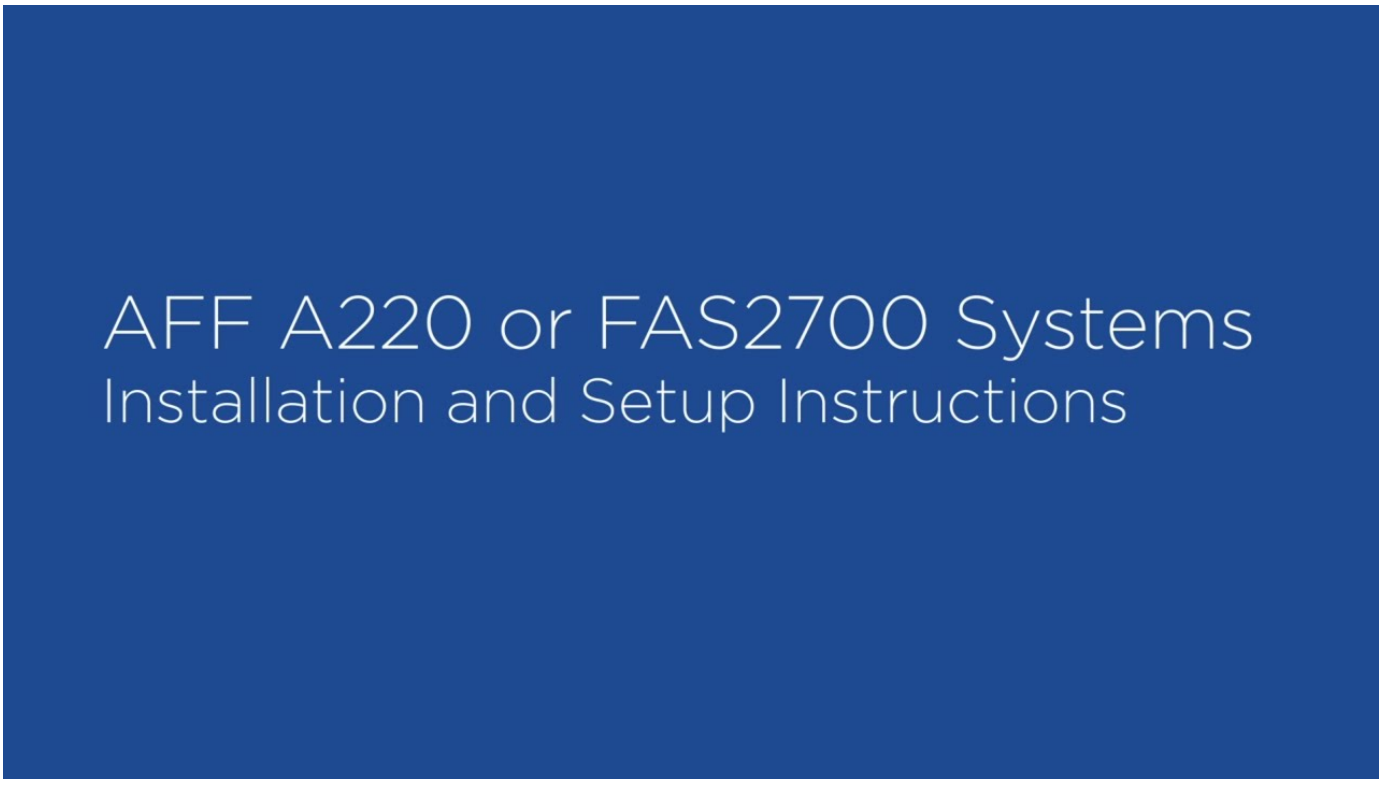

Detailed guide - AFF A220

This page gives detailed step-by-step instructions for installing a typical NetApp system.

Step 1: Prepare for installation

To install your AFF A220 system, you need to create an account on the NetApp Support Site, register your system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser
- A laptop or console with an RJ-45 connection and access to a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register your system.

      NetApp Product Registration

4. Download and install Config Advisor on your laptop.

   NetApp Downloads: Config Advisor

5. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
10 GbE cable (order dependent)	X6566B-05-R6 (112-00297), 0.5m X6566B-2-R6 (112-00299), 2m		Cluster interconnect network

Type of cable…	Part number and length	Connector type	For…
10 GbE cable (order dependent)	Part number X6566B-2-R6 (112-00299), 2m or X6566B-3-R6 (112-00300), 3m X6566B-5-R6 (112-00301), 5m		Data
Optical network cables (order dependent)	X6553-R6 (112-00188), 2m X6536-R6 (112-00090), 5m X6554-R6(112-00189), 15m		FC host network
Cat 6, RJ-45 (order dependent)	Part numbers X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network and Ethernet data
Storage (order dependent)	Part number X66030A (112-00435), 0.5m X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		Storage
Micro-USB console cable	Not applicable		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	Not applicable		Powering up the system

6. Download and complete the *Cluster configuration worksheet*.

   Cluster Configuration Worksheet

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

3. Attach cable management devices (as shown).

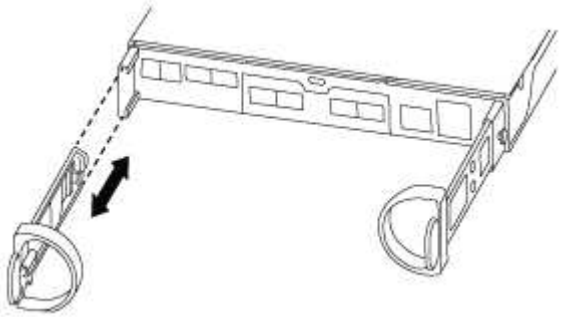

4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Cable a two-node switchless cluster, unified network configuration

Management network, UTA2 data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

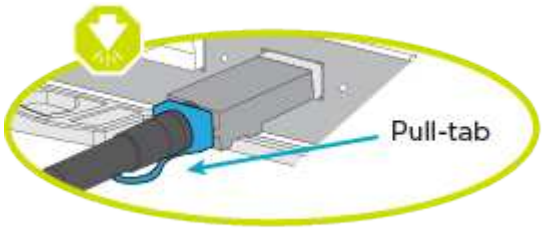

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and to the switches:

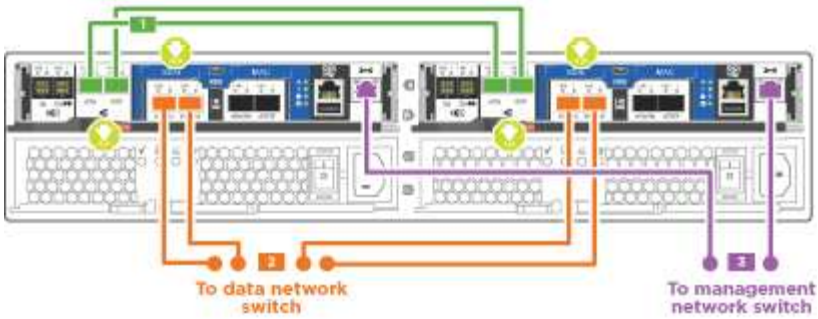

Step	Perform on each controller
1	Cable the cluster interconnect ports to each other with the cluster interconnect cable: • e0a to e0a • e0b to e0b 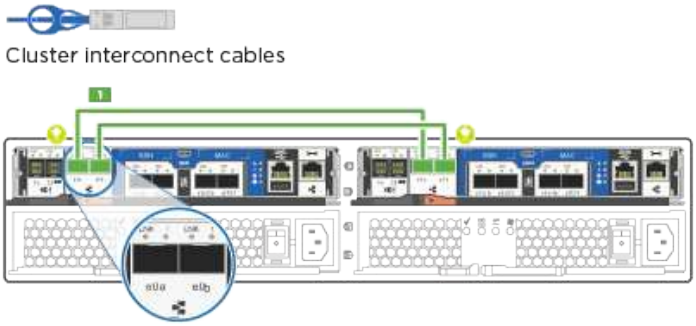

Step	Perform on each controller
2	Use one of the following cable types to cable the UTA2 data ports to your host network: An FC host • 0c and 0d • **or** 0e and 0f A 10GbE • e0c and e0d • **or** e0e and e0f You can connect one port pair as CNA and one port pair as FC, or you can connect both port pairs as CNA or both port pairs as FC. 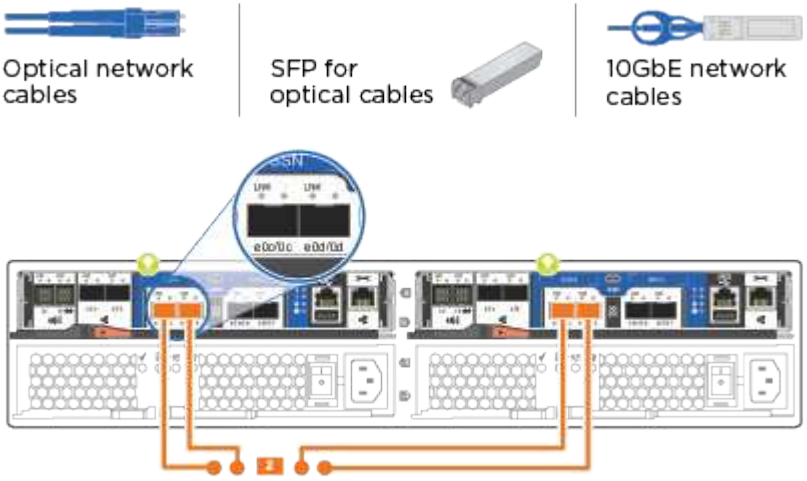
3	Cable the e0M ports to the management network switches with the RJ45 cables: 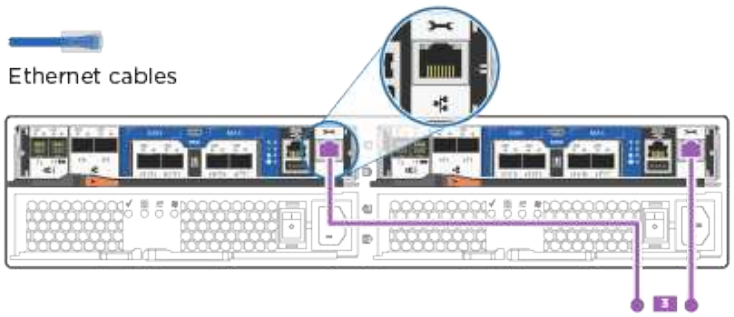
	DO NOT plug in the power cords at this point.

2. To cable your storage, see Step 4: Cable controllers to drive shelves

Option 2: Cable a switched cluster, unified network configuration

Management network, UTA2 data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled to the cluster interconnect switches.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

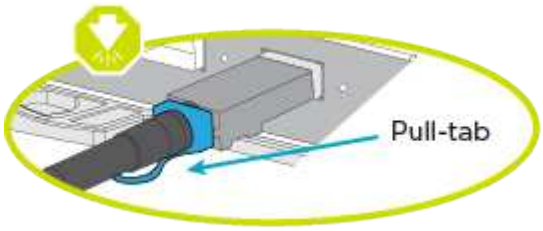

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and the switches:

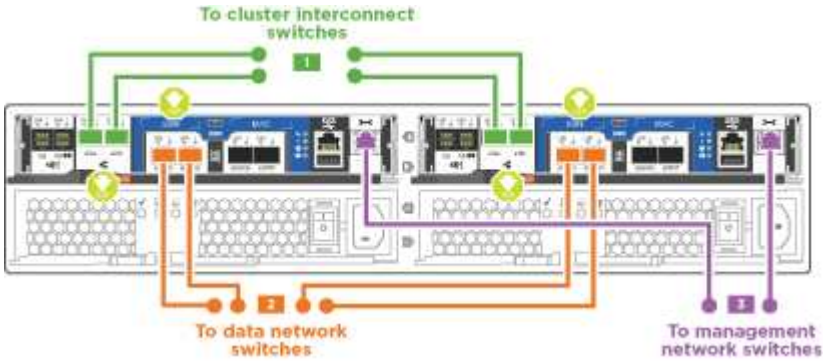

Step	Perform on each controller module
1	Cable e0a and e0b to the cluster interconnect switches with the cluster interconnect cable: 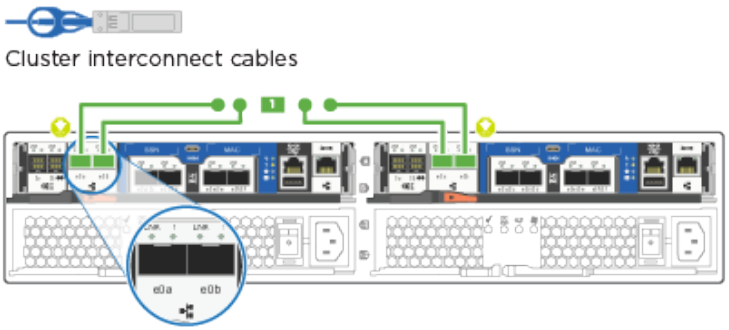
2	Use one of the following cable types to cable the UTA2 data ports to your host network: An FC host • 0c and 0d • **or** 0e and 0f A 10GbE • e0c and e0d • **or** e0e and e0f You can connect one port pair as CNA and one port pair as FC, or you can connect both port pairs as CNA or both port pairs as FC. 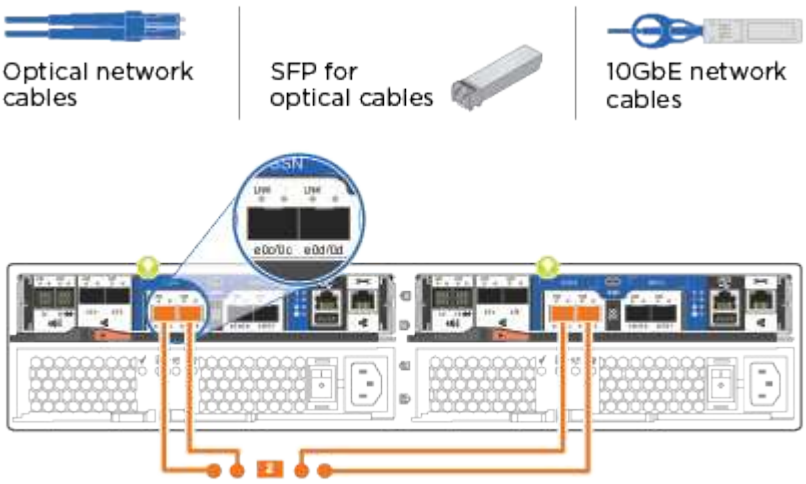

Step	Perform on each controller module
3	Cable the e0M ports to the management network switches with the RJ45 cables: 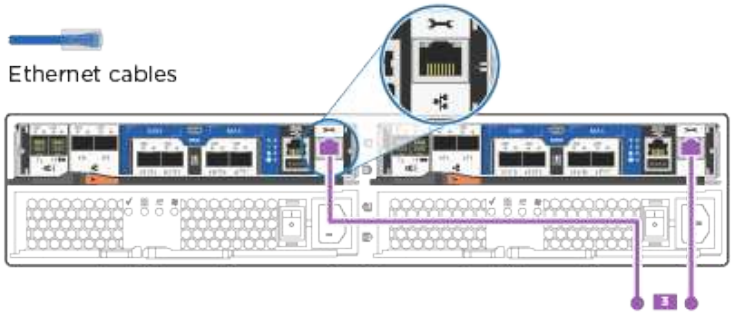
	DO NOT plug in the power cords at this point.

2. To cable your storage, see Step 4: Cable controllers to drive shelves

Option 3: Cable a two-node switchless cluster, Ethernet network configuration

Management network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

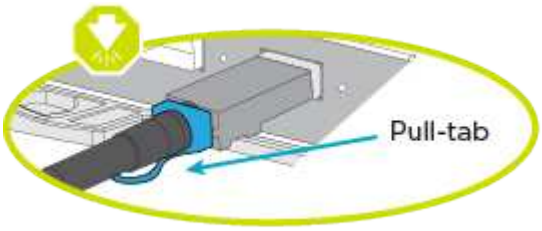

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and to the switches:

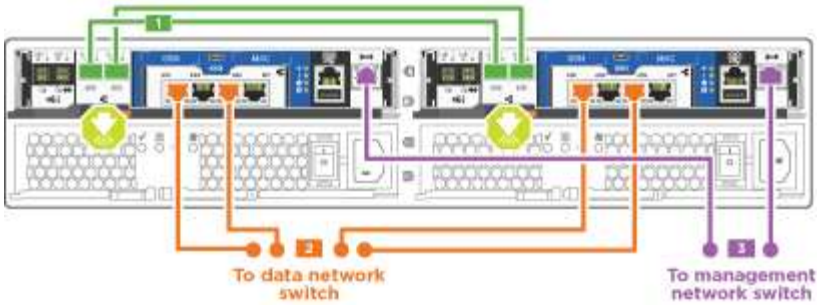

Step	Perform on each controller
1	Cable the cluster interconnect ports to each other with the cluster interconnect cable: • e0a to e0a • e0b to e0b 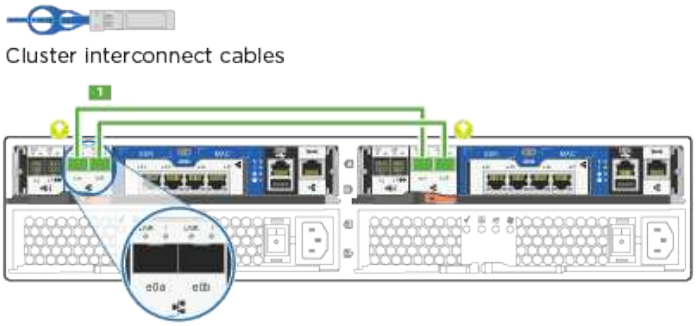
2	Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network: 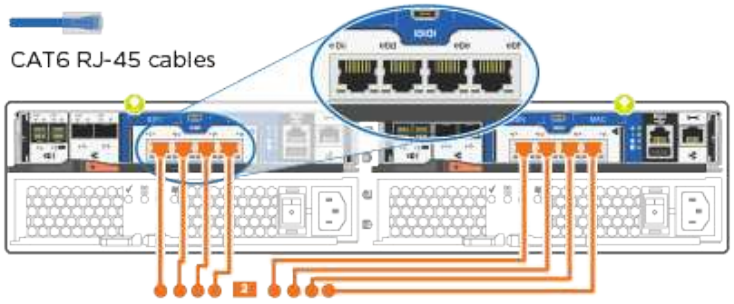

Step	Perform on each controller
3	Cable the e0M ports to the management network switches with the RJ45 cables: 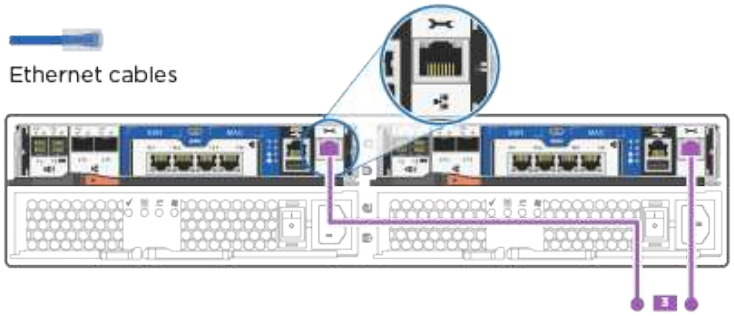
	DO NOT plug in the power cords at this point.

2. To cable your storage, see Step 4: Cable controllers to drive shelves

Option 4: Cable a switched cluster, Ethernet network configuration

Management network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled to the cluster interconnect switches.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

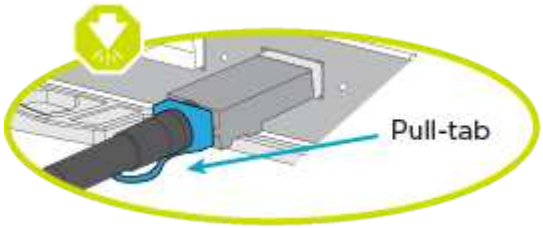

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and the switches:

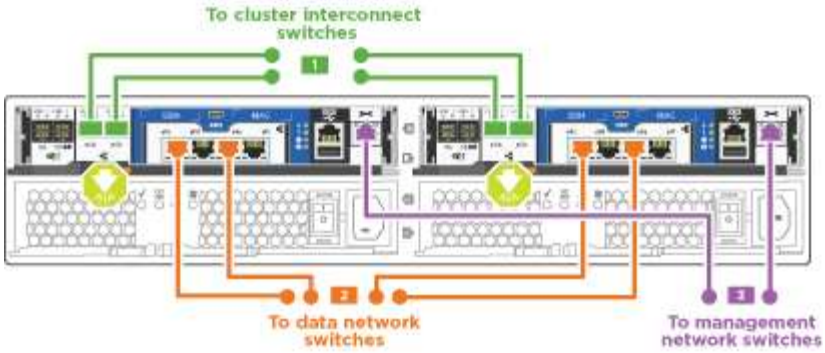

Step	Perform on each controller module
1	Cable e0a and e0b to the cluster interconnect switches with the cluster interconnect cable: 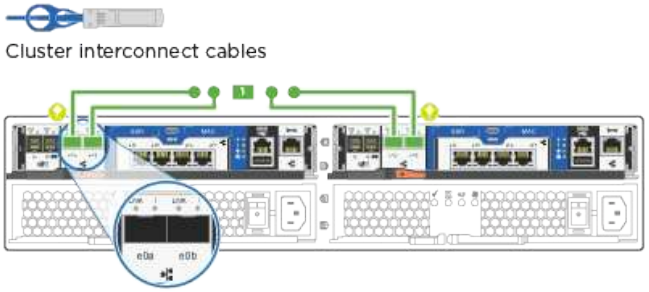
2	Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network: 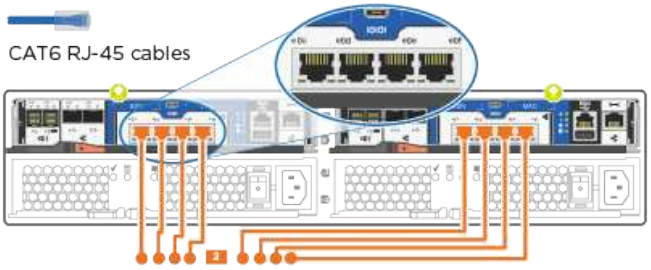

Step	Perform on each controller module
3	Cable the e0M ports to the management network switches with the RJ45 cables: 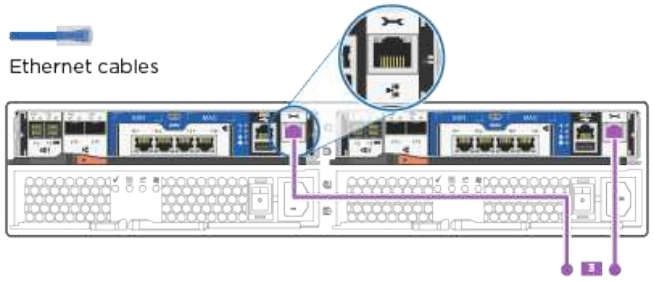
	DO NOT plug in the power cords at this point.

2. To cable your storage, see Step 4: Cable controllers to drive shelves

Step 4: Cable controllers to drive shelves

You must cable the controllers to your shelves using the onboard storage ports. NetApp recommends MP-HA cabling for systems with external storage. If you have a SAS tape drive, you can use single-path cabling. If you have no external shelves, MP-HA cabling to internal drives is optional (not shown) if the SAS cables are ordered with the system.

Option 1: Cable storage on an HA pair with external drive shelves

You must cable the shelf-to-shelf connections, and then cable both controllers to the drive shelves.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

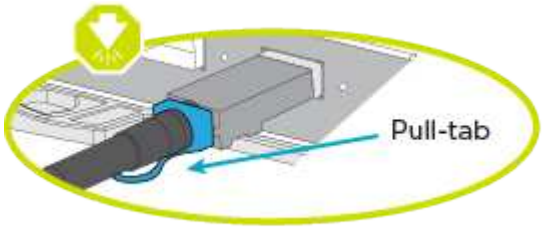

Steps

1. Cable the HA pair with external drive shelves:

> The example uses DS224C. Cabling is similar with other supported drive shelves.

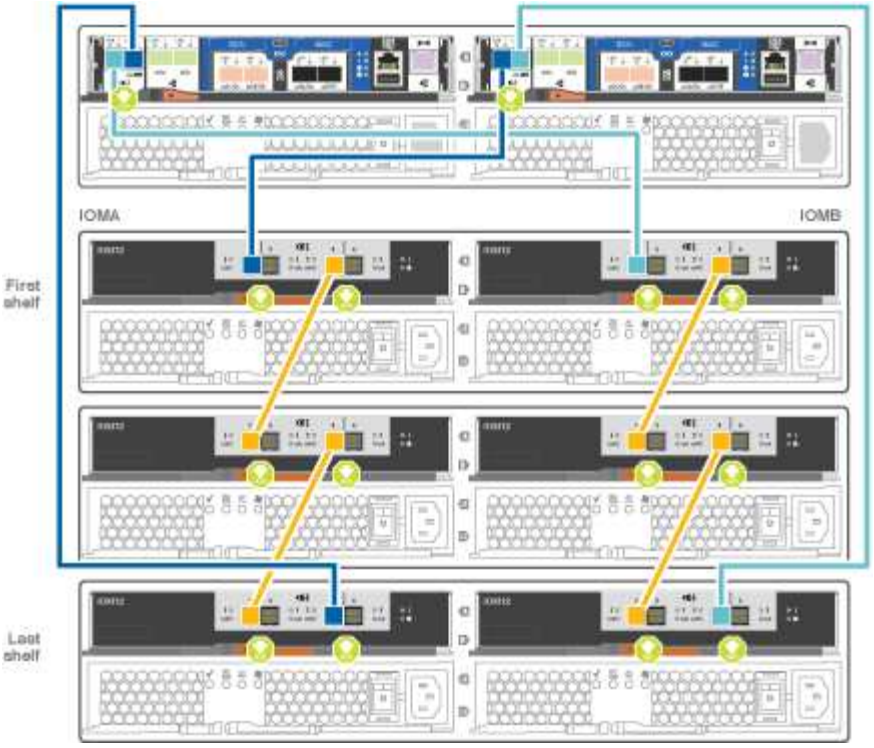

Step	Perform on each controller
1	Cable the shelf-to-shelf ports. • Port 3 on IOM A to port 1 on the IOM A on the shelf directly below. • Port 3 on IOM B to port 1 on the IOM B on the shelf directly below.  mini-SAS HD to mini-SAS HD cables
2	Connect each node to IOM A in the stack. • Controller 1 port 0b to IOM A port 3 on last drive shelf in the stack. • Controller 2 port 0a to IOM A port 1 on the first drive shelf in the stack.  mini-SAS HD to mini-SAS HD cables
3	Connect each node to IOM B in the stack • Controller 1 port 0a to IOM B port 1 on first drive shelf in the stack. • Controller 2 port 0b to IOM B port 3 on the last drive shelf in the stack.  mini-SAS HD to mini-SAS HD cables

If you have more than one drive shelf stack, see the *Installation and Cabling Guide* for your drive shelf type.

2. To complete setting up your system, see Step 5: Complete system setup and configuration

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Complete system setup if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Use the following animation to set one or more drive shelf IDs

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Turn on the power switches to both nodes.

Initial booting may take up to eight minutes.

4. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

5. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

6. Select an ONTAP icon listed to discover:

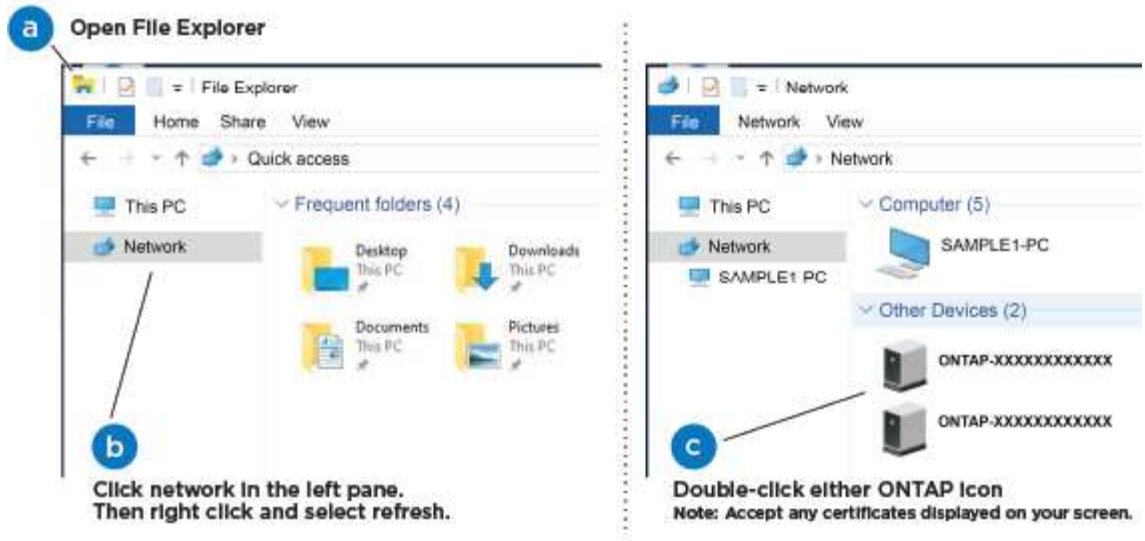

a. Open File Explorer.

b. Click **Network** in the left pane and right-click and select **refresh**.

c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

   XXXXX is the system serial number for the target node.

   System Manager opens.

7. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

   ONTAP Configuration Guide

8. Verify the health of your system by running Config Advisor.

9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:

   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your system.

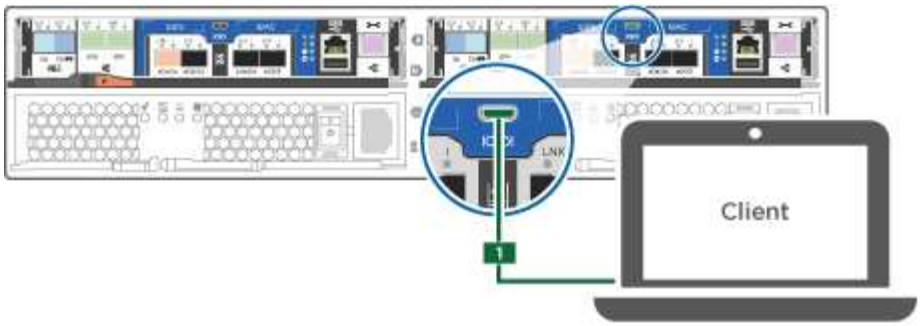

c. Connect the laptop or console to the switch on the management subnet.

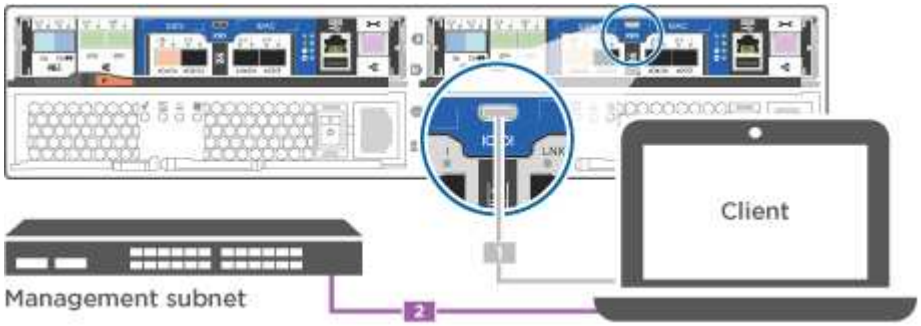

d. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Use the following animation to set one or more drive shelf IDs:

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
4. Turn on the power switches to both nodes.

Initial booting may take up to eight minutes.

5. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.

If the management network has DHCP…	Then…
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

6. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the *NetApp ONTAP Configuration guide*.

      ONTAP Configuration Guide

7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain AFF A220 hardware

For the AFF A220 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

NVEM Battery

A battery is included with a controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF A220

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct node:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy node* is the HA partner of the impaired node.

Check encryption key support and status - AFF A220

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the impaired controller - AFF A220

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller displays…	Then…
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - AFF A220

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

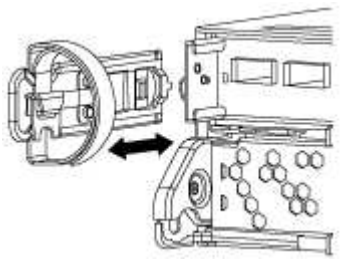

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

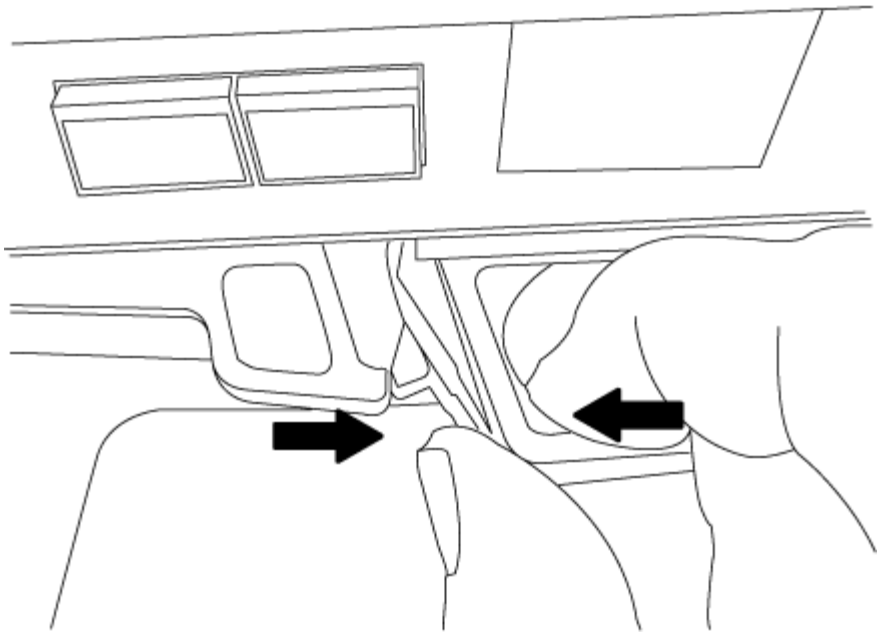

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

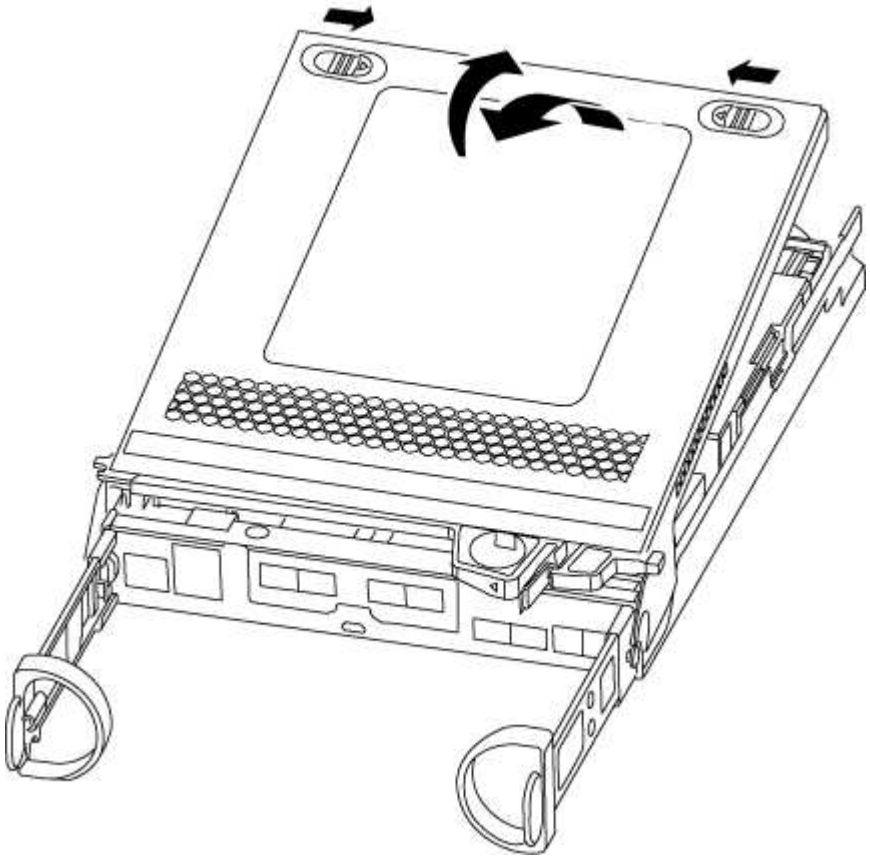

Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the boot media using the following illustration or the FRU map on the controller module:

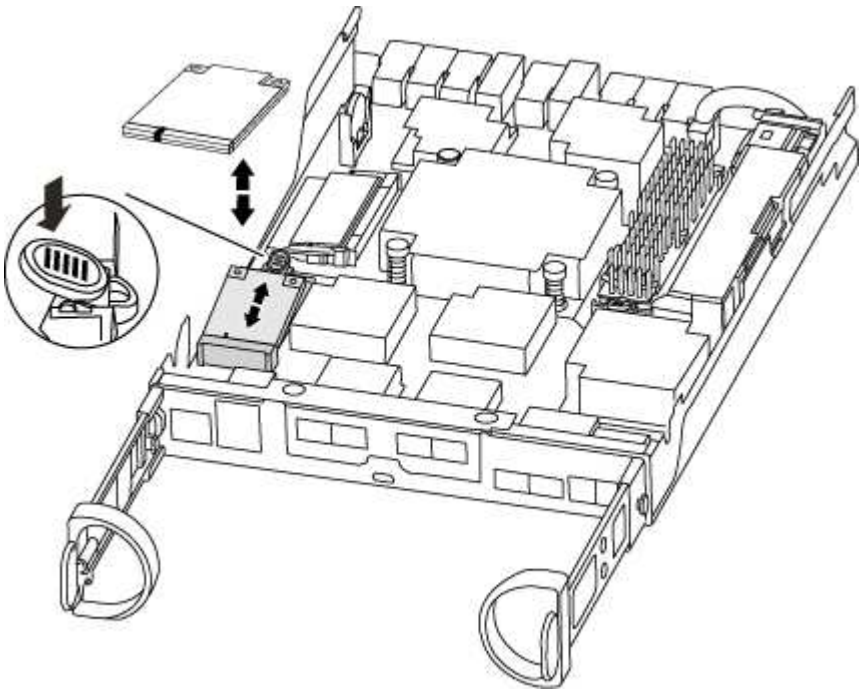

3. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

4. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
5. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

6. Push the boot media down to engage the locking button on the boot media housing.
7. Close the controller module cover.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

   The controller begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

6. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

   The system begins to boot and stops at the LOADER prompt.

7. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

     

     The target port you configure is the target port you use to communicate with the impaired controller from the healthy controller during var file system restore with a network connection. You can also use the e0M port in this command.

   - If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
     - filer_addr is the IP address of the storage system.
     - netmask is the network mask of the management network that is connected to the HA partner.
     - gateway is the gateway for the network.
     - dns_addr is the IP address of a name server on your network.
     - dns_domain is the Domain Name System (DNS) domain name.

       If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

Boot the recovery image - AFF A220

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:

- If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed
- If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.
8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore encryption - AFF A220

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.

- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*******************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*******************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    `security key-manager key query -restored false.`

    ```
    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    `storage failover giveback -fromnode local`

13. Restore automatic giveback, if you disabled it, by entering the following command.

    `storage failover modify -node local -auto-giveback true`

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    `system node autosupport invoke -node * -type all -message MAINT=END`

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*********************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*********************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF A220

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF A220

To replace the chassis, you must move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving all drives and controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-controller cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A220

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous

step to monitor progress.

If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

```
Warning: Are you sure you want to halt node <node_name>? {y|n}:
```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Controller is in a MetroCluster configuration

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Move and replace hardware - AFF A220

Move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Move a power supply

Moving out a power supply when replacing a chassis involves turning off, disconnecting, and removing the power supply from the old chassis and installing and connecting it on the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.
4. Use the cam handle to slide the power supply out of the system.

   

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

Step 2: Remove the controller module

Remove the controller module or modules from the old chassis.

1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

2. Remove and set aside the cable management devices from the left and right sides of the controller module.

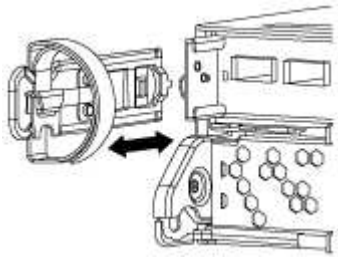

3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

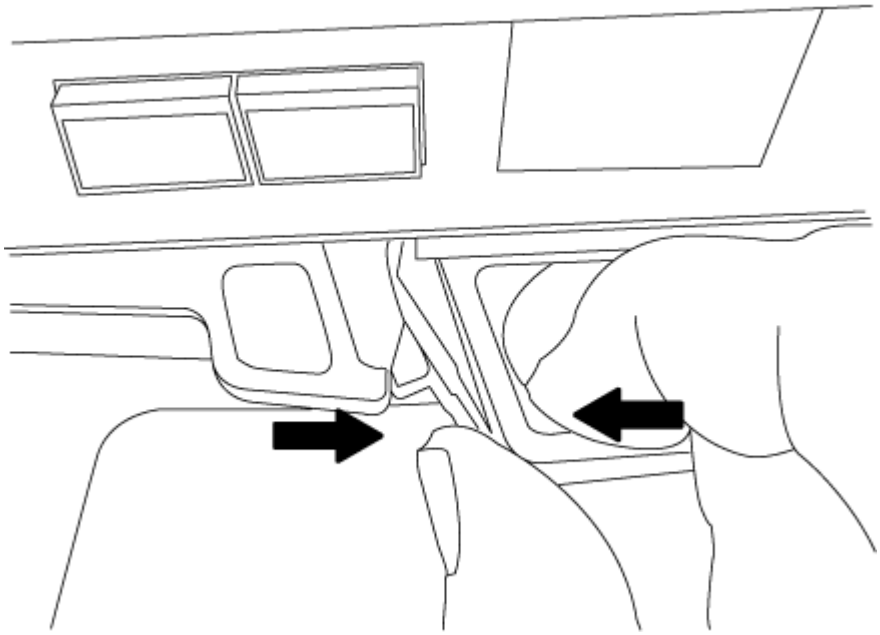

4. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 3: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new

chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.
5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It click when it is secure.
6. Repeat the process for the remaining drives in the system.

Step 4: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller

After you install the controller module and any other components into the new chassis, boot it the system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Repeat the preceding steps if there is a second controller to install in the new chassis.
4. Complete the installation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Repeat the preceding steps for the second controller module in the new chassis.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reinstall the blanking panel and then go to the next step.

5. Connect the power supplies to different power sources, and then turn them on.
6. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - AFF A220

You must verify the HA state of the chassis, switch back aggregates, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mcc-2n`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. The next step depends on your system configuration.
5. Reboot the system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - AFF A220

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- This procedure includes steps for automatically or manually reassigning drives to the *replacement* controller, depending on your system’s configuration.

  You should perform the drive reassignment as directed in the procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - AFF A220

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - AFF A220

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove controller module

To replace the controller module, you must first remove the old controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

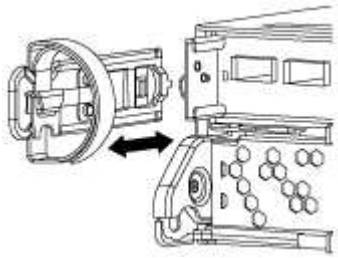

4. If you left the SFP modules in the system after removing the cables, move them to the new controller module.
5. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

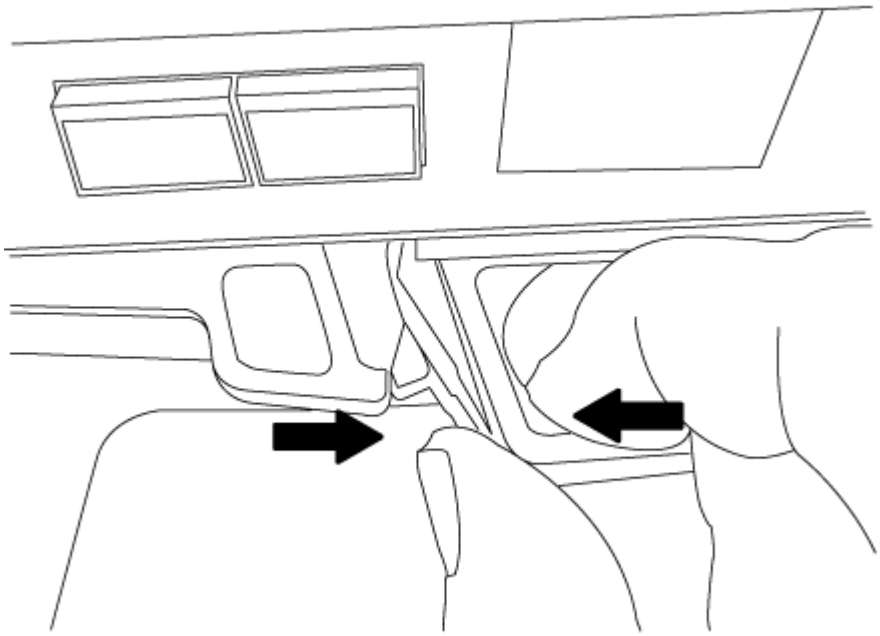

6. Turn the controller module over and place it on a flat, stable surface.
7. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

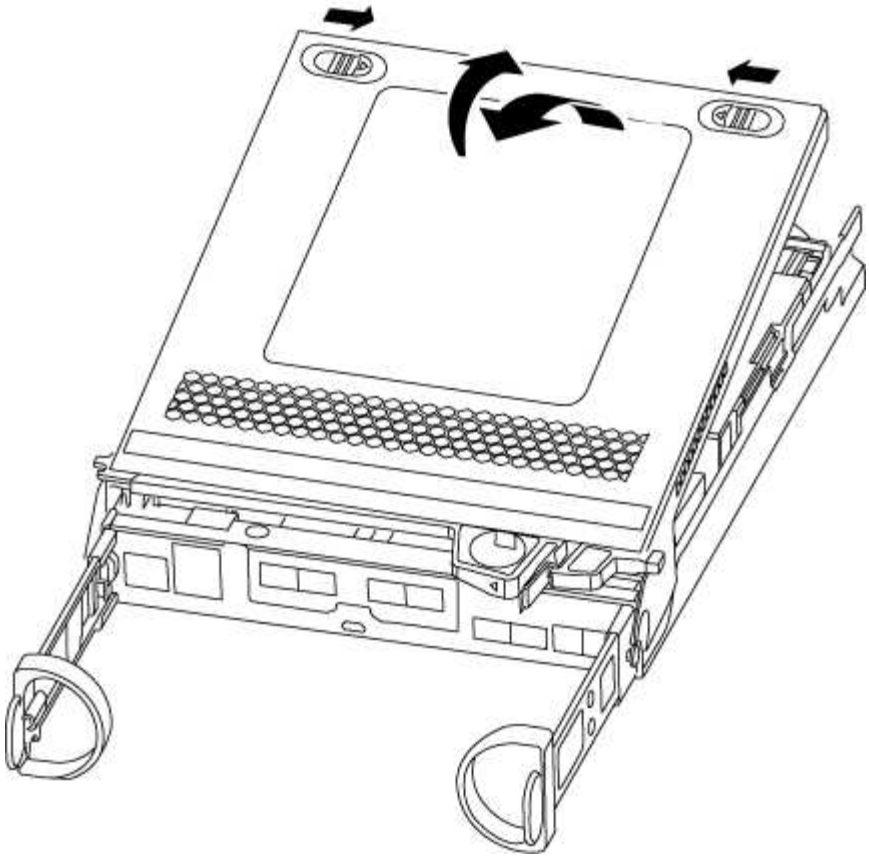

Step 2: Move the NVMEM battery

To move the NVMEM battery from the old controller module to the new controller module, you must perform a specific sequence of steps.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Locate the NVMEM battery in the controller module.

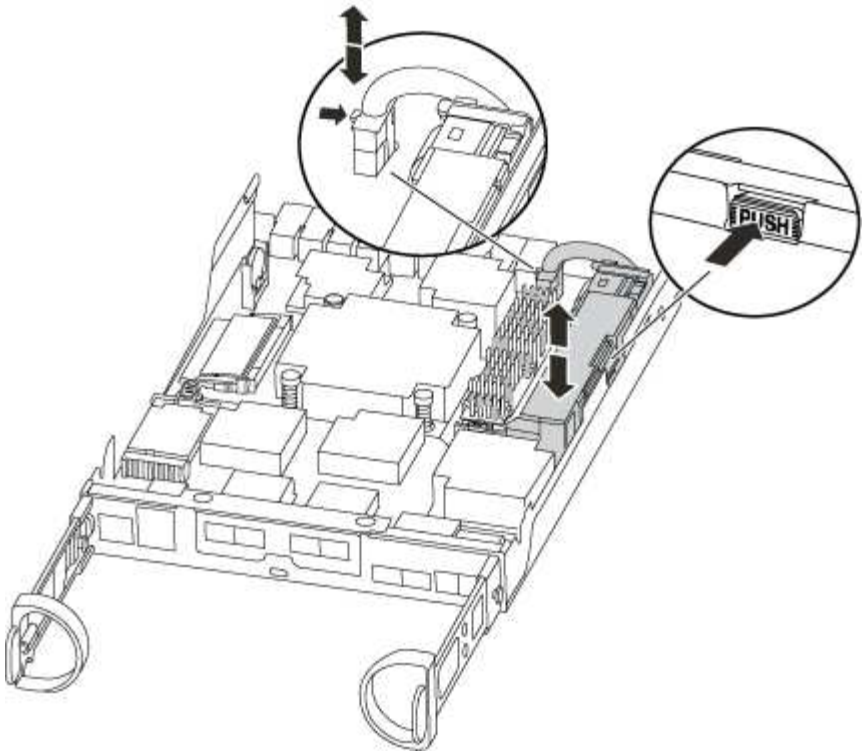

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Loop the battery cable around the cable channel on the side of the battery holder.
7. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
8. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.

Step 3: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller module and insert it in the new controller module.

1. Locate the boot media using the following illustration or the FRU map on the controller module:

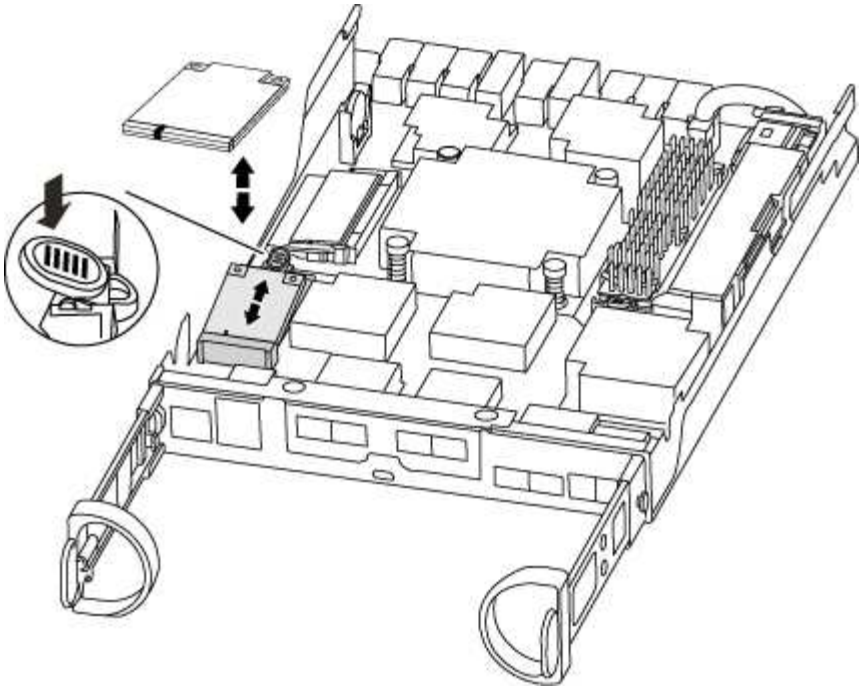

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 4: Move the DIMMs

To move the DIMMs, you must follow the directions to locate and move them from the old controller module into the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

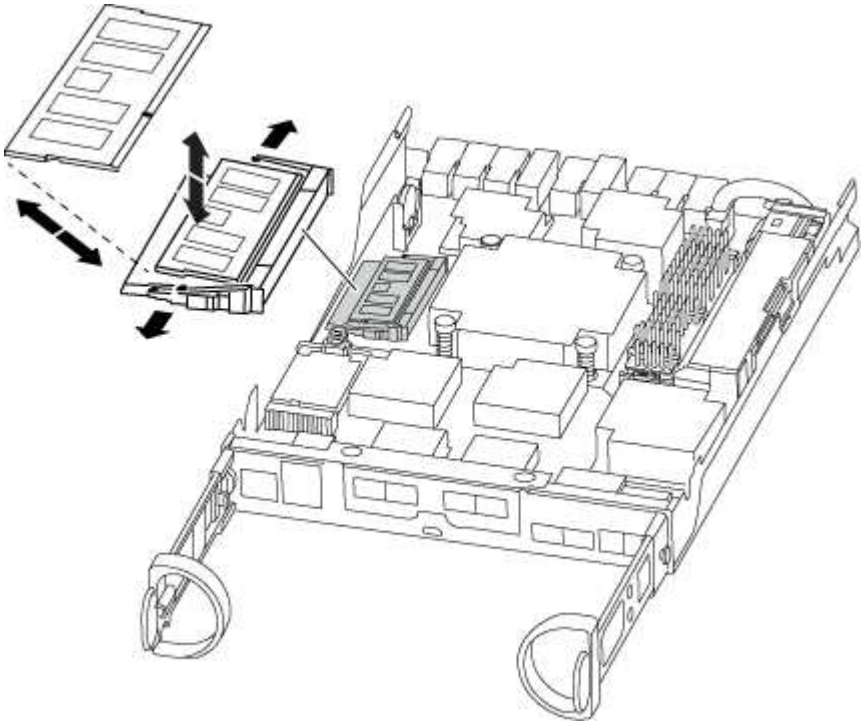

4. Repeat these steps to remove additional DIMMs as needed.
5. Verify that the NVMEM battery is not plugged into the new controller module.
6. Locate the slot where you are installing the DIMM.
7. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Repeat these steps for the remaining DIMMs.
9. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

   Make sure that the plug locks down onto the controller module.

Step 5: Move a caching module, if present

If your AFF A220 or FAS2700 system has a caching module, you need to move the caching module from the old controller module to the replacement controller module. The caching module is referred to as the “M.2 PCIe card” on the controller module label.

You must have the new controller module ready so that you can move the caching module directly from the old controller module to the corresponding slot in the new one. All other components in the storage system must be functioning properly; if not, you must contact technical support.

1. Locate the caching module at the rear of the controller module and remove it.

a. Press the release tab.

b. Remove the heatsink.

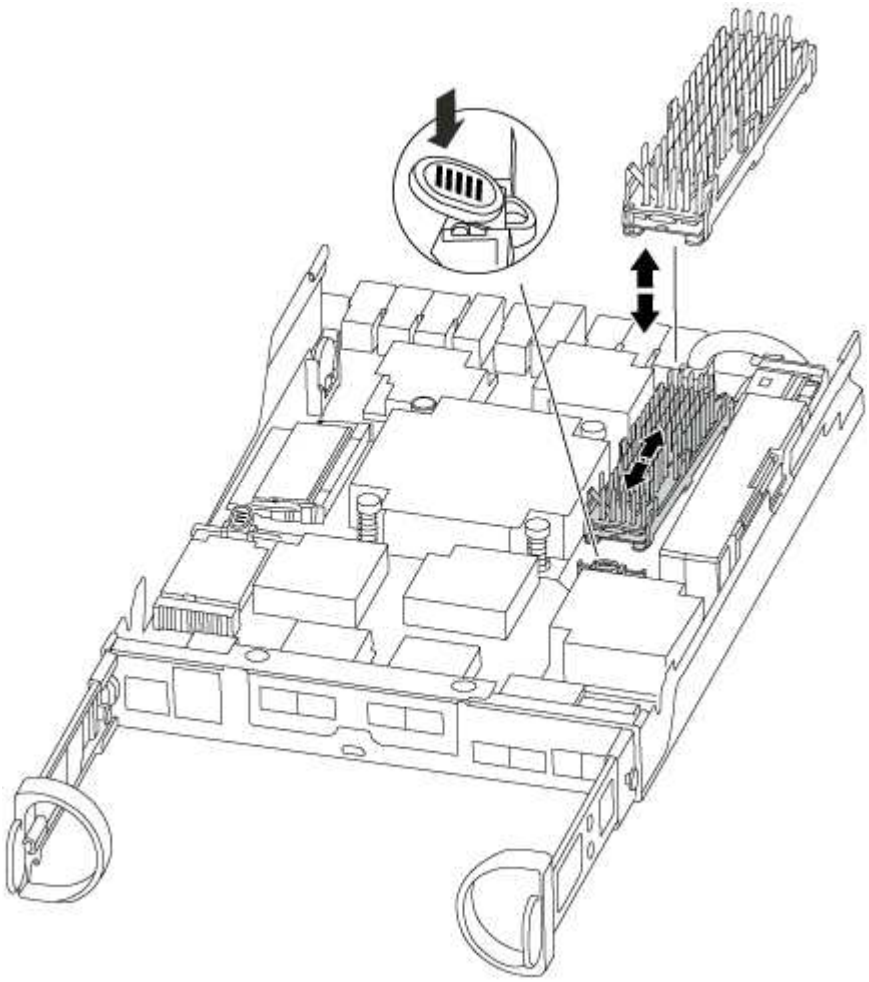

2. Gently pull the caching module straight out of the housing.
3. Move the caching module to the new controller module, and then align the edges of the caching module with the socket housing and gently push it into the socket.
4. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

5. Reseat and push the heatsink down to engage the locking button on the caching module housing.
6. Close the controller module cover, as needed.

Step 6: Install the controller

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Interrupt the boot process **only** after determining the correct timing: You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. Only press `Ctrl-C` when you see the message `Press Ctrl-C for Boot Menu`. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. Select the option to boot to Maintenance mode from the displayed menu.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process. e. Interrupt the boot process **only** after determining the correct timing: You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. Only press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. f. From the boot menu, select the option for Maintenance mode.

Important: During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
  You can safely respond `y` to these prompts.

Restore and verify the system configuration - AFF A220

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the controller's HA state

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

- `ha`
- `mcc`
- `mcc-2n`
- `mccip`
- `non-ha`

b. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF A220

To complete the replacement procedure and restore your system to full operation, you must recable the storage, confirm disk reassignment, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks.

You must use the correct procedure for your configuration:

Controller redundancy	Then use this procedure…
HA pair	Option 1: Verify the system ID change on an HA system
Stand-alone	Option 2: Manually reassign the system ID on a stand-alone system in ONTAP
Two-node MetroCluster configuration	Option 3: Manually reassign the system ID on systems in a two-node MetroCluster configuration

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node` *`local-node-name`* `partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node` *`local-node-name`* `partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

- Restore onboard key management encryption keys
- Restore external key management encryption keys

6. Give back the controller:

   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Option 2: Manually reassign the system ID on a stand-alone system in ONTAP

In a stand-alone system, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems that are in a stand-alone configuration.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by pressing Ctrl-C, and then select the option to boot to Maintenance mode from the displayed menu.
2. You must enter `Y` when prompted to override the system ID due to a system ID mismatch.
3. View the system IDs: `disk show -a`
4. You should make a note of the old system ID, which is displayed as part of the disk owner column.

   The following example shows the old system ID of 118073209:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                  POOL   SERIAL NUMBER  HOME
--------    -------------          -----  -------------  -------------
disk_name    system-1  (118073209)  Pool0  J8XJE9LC       system-1
(118073209)
disk_name    system-1  (118073209)  Pool0  J8Y478RC       system-1
(118073209)
.
.
.
```

5. Reassign disk ownership by using the system ID information obtained from the disk show command: `disk reassign -s old system ID disk reassign -s 118073209`
6. Verify that the disks were assigned correctly: `disk show -a`

   The disks belonging to the replacement node should show the new system ID. The following example now show the disks owned by system-1 the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                  POOL   SERIAL NUMBER  HOME
--------    -------------          -----  -------------  -------------
disk_name    system-1  (118065481)  Pool0  J8Y0TDZC       system-1
(118065481)
disk_name    system-1  (118065481)  Pool0  J8Y0TDZC       system-1
(118065481)
.
.
.
```

7. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:

- Restore onboard key management encryption keys
- Restore external key management encryption keys

8. Boot the node: `boot_ontap`

Option 3: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `metrocluster node show -fields node-systemid,dr-partner-systemid`

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

   ```
   dr-group-id cluster         node                 node-systemid dr-
   partner-systemid
    ----------- --------------------- ------------------- -------------
   -------------------
    1           Cluster_A             Node_A_1             536872914
   118073209
    1           Cluster_B             Node_B_1             118073209
   536872914
    2 entries were displayed.
   ```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

   ```
   Local System ID: 118065481
       ...
       ...
   ```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                  POOL   SERIAL NUMBER  HOME
-------    -------------          -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Verify that the coredumps are saved: `system node run -node local-node-name partner savecore`

      If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node local-node-name partner savecore -s command`.</info>.

   c. Return to the admin privilege level: `set -privilege admin`

7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`

8. Boot the *replacement* node: `boot_ontap`

9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`

10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

       You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

    b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`
    c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - AFF A220

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.

4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A220

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

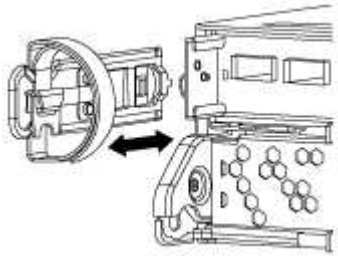

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

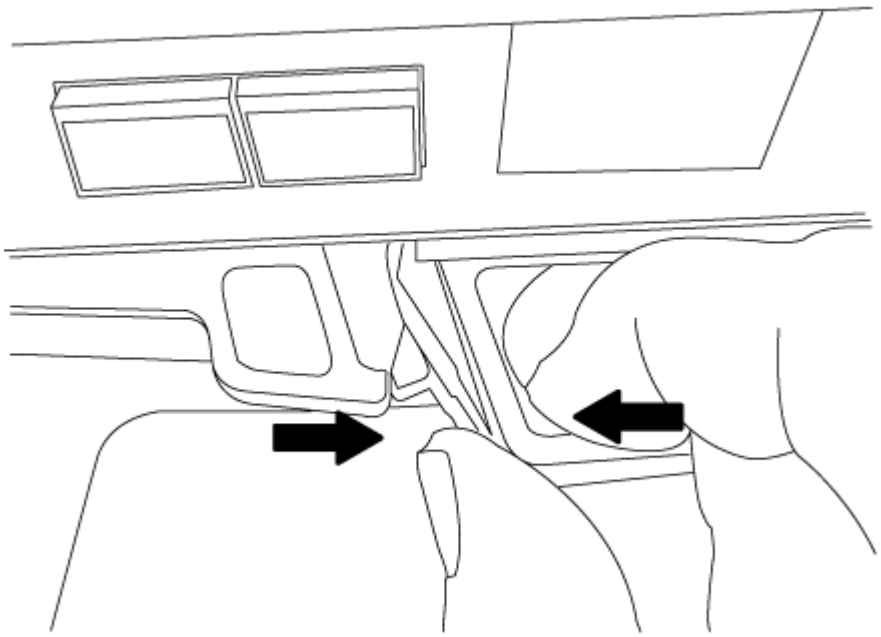

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

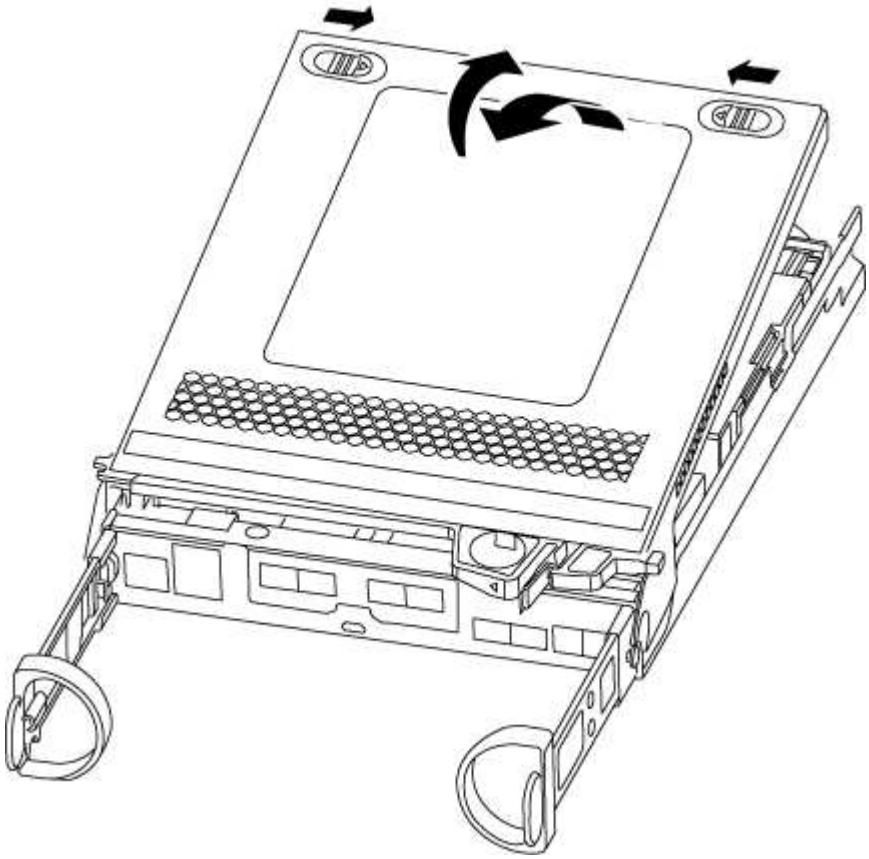

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

If you are replacing a DIMM, you need to remove it after you have unplugged the NVMEM battery from the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED on the back of controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

   

3. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
4. If the NVMEM LED is flashing, there is data in the NVMEM and you must disconnect the battery to clear the memory:
   a. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

   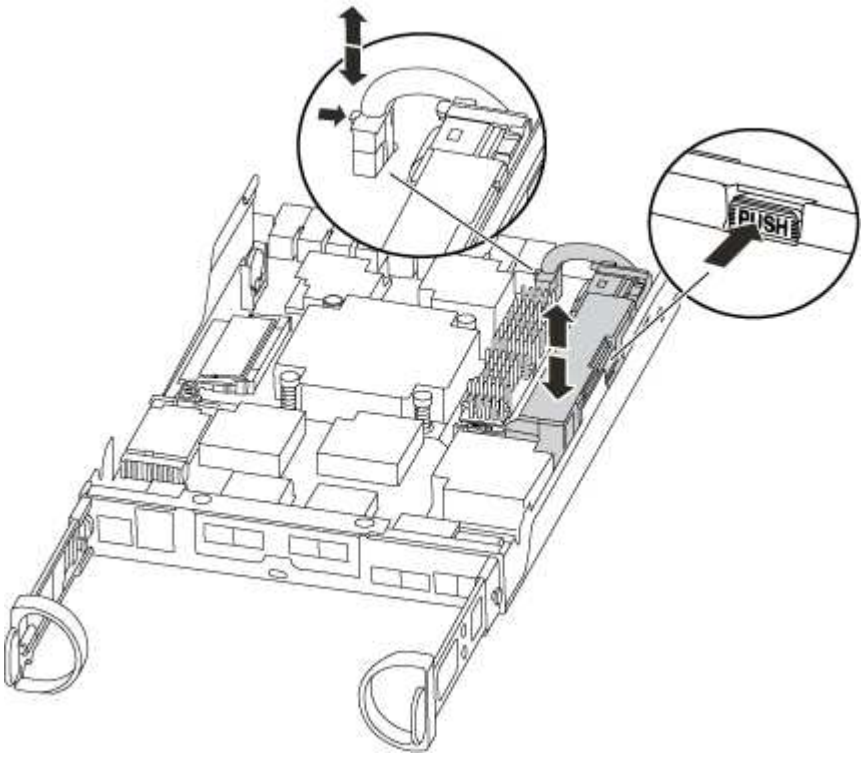

   b. Confirm that the NVMEM LED is no longer lit.
   c. Reconnect the battery connector.

5. Return to Step 3: Replace the DIMMs in this procedure to recheck the NVMEM LED.
6. Locate the DIMMs on your controller module.
7. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
8. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   The number and placement of system DIMMs depends on the model of your system.

   The following illustration shows the location of system DIMMs:

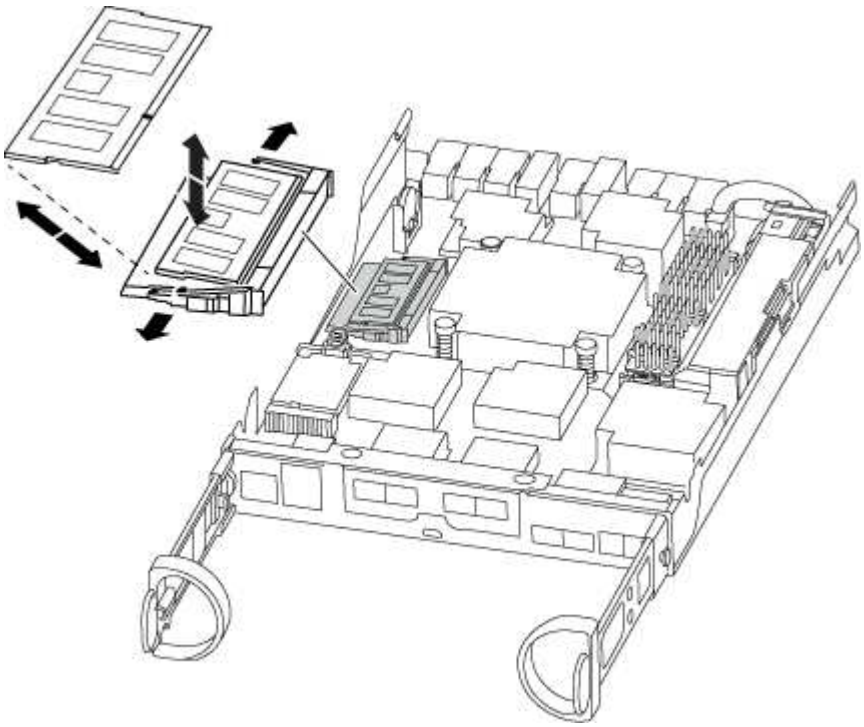

9. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

10. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

    The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

    Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

11. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
12. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

Make sure that the plug locks down onto the controller module.

13. Close the controller module cover.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - AFF A220

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace the NVMEM battery - AFF A220

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

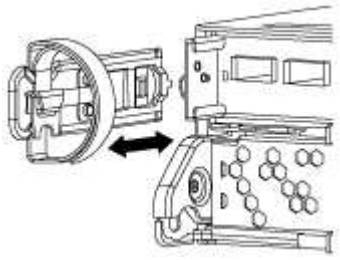

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

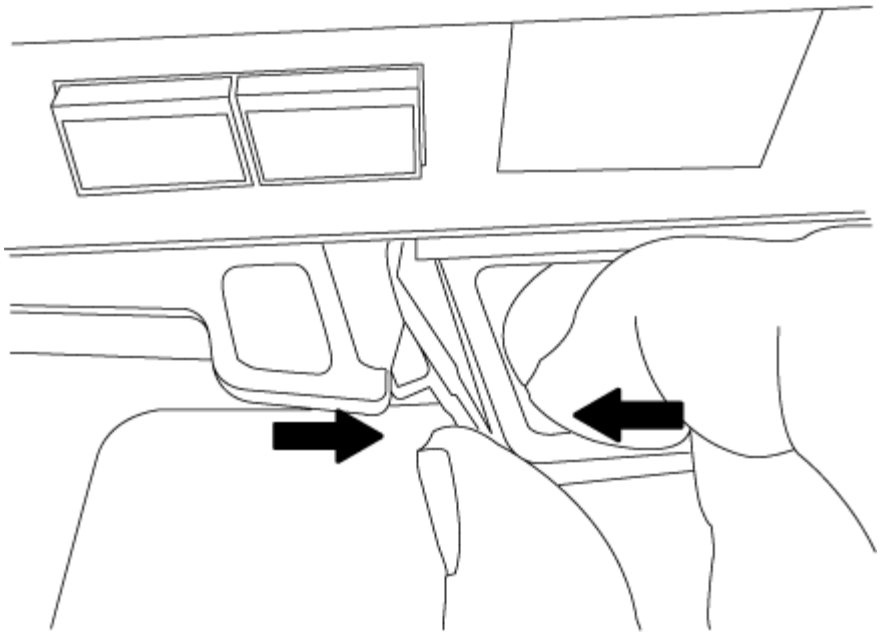

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery in your system, you must remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

3. Locate the NVMEM battery in the controller module.

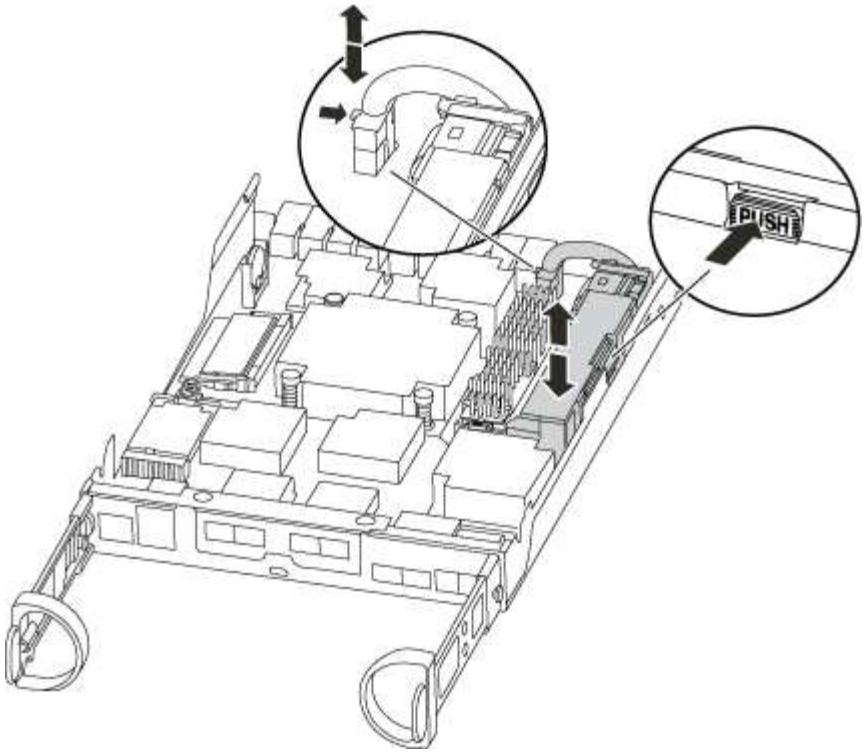

4. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
5. Remove the battery from the controller module and set it aside.
6. Remove the replacement battery from its package.
7. Loop the battery cable around the cable channel on the side of the battery holder.
8. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
9. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
10. Plug the battery plug back into the controller module.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber

optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - AFF A220

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> Cooling is integrated with the power supply, so you must replace the power supply within two minutes of removal to prevent overheating due to reduced airflow. Because the chassis provides a shared cooling configuration for the two HA nodes, a delay longer than two minutes will shut down all controller modules in the chassis. If both controller modules do shut down, make sure that both power supplies are inserted, turn both off for 30 seconds, and then turn both on.

- Power supplies are auto-ranging.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.

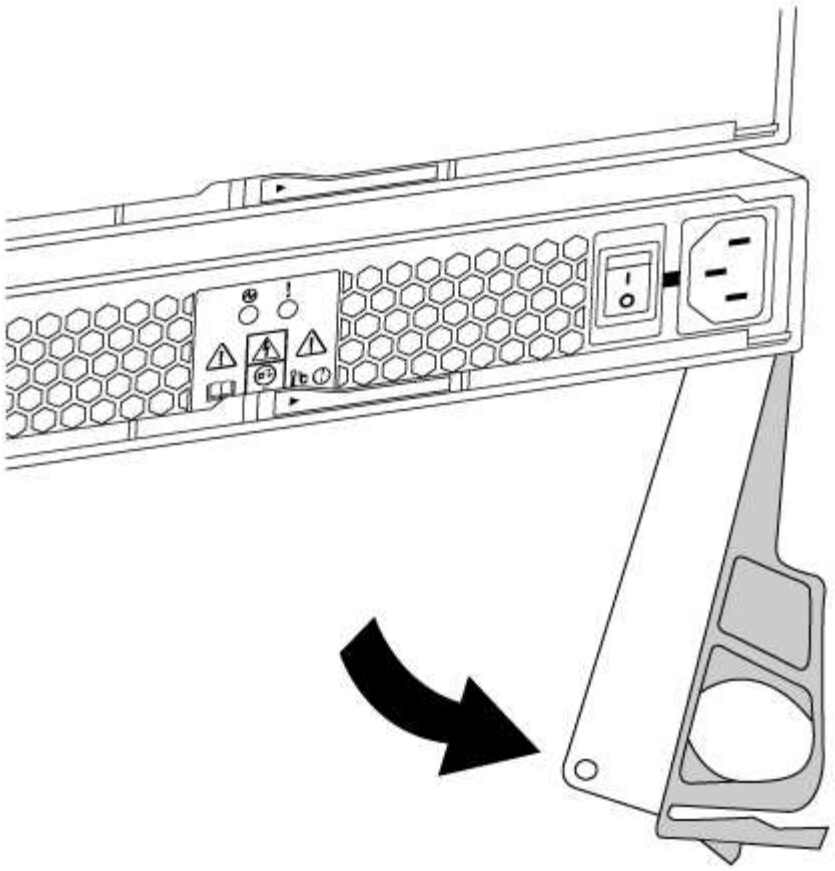

5. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.
7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
9. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.

10. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

    The power supply LEDs are lit when the power supply comes online.

11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A220

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

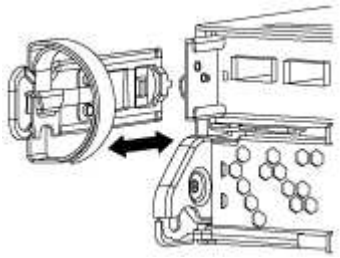

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

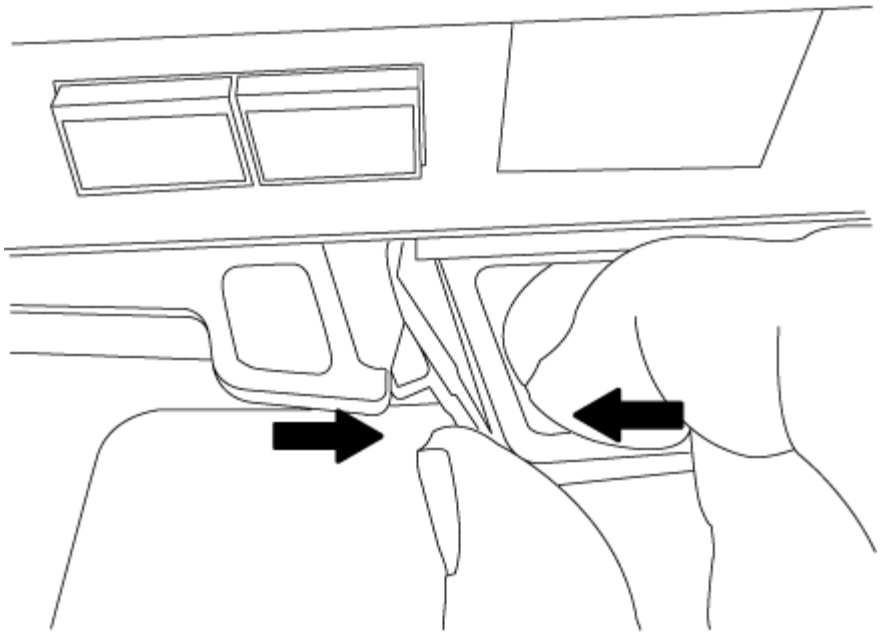

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

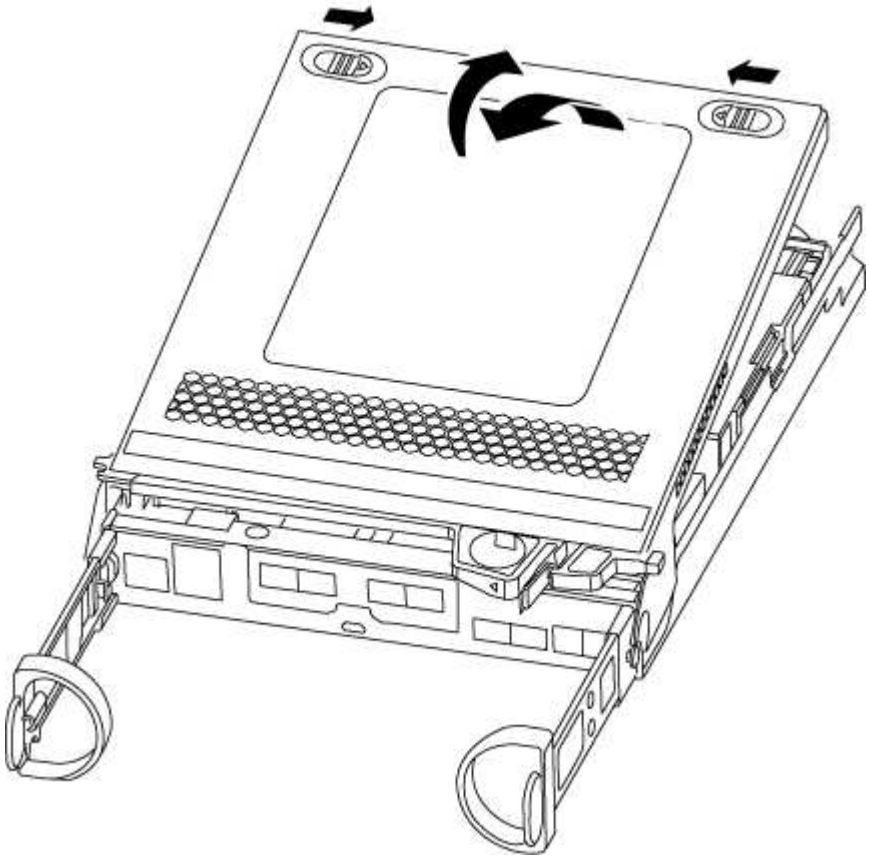

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

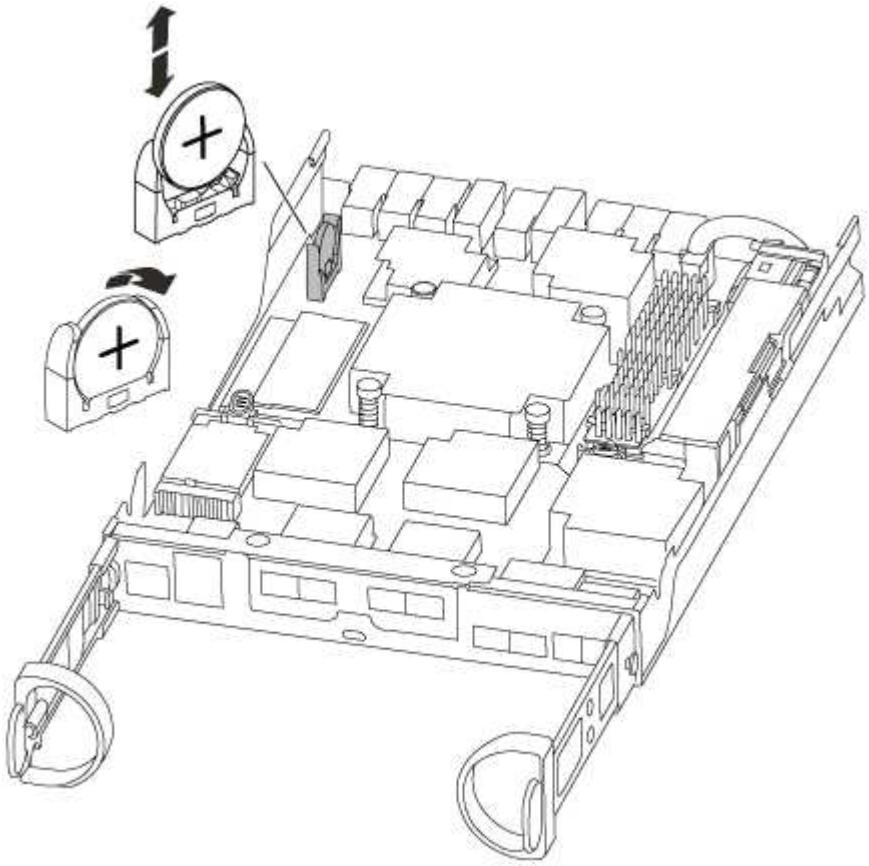

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and set time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF A300 systems

Install and setup

Cluster configuration worksheet - AFF A300

You can use the worksheet to gather and record your site-specific IP addresses and other information required when configuring an ONTAP cluster.

Cluster Configuration Worksheet

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

For MetroCluster configurations, see either:

- Install MetroCluster IP configuration
- Install MetroCluster Fabric-Attached configuration

Installation and setup PDF poster - AFF A300

You can use the PDF poster to install and set up your new system. The PDF poster provides step-by-step instructions with live links to additional content.

AFF A300 Installation and Setup Instructions

Maintain

Maintain AFF A300 hardware

For the AFF A300 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Fan

The fan cools the controller.

NVDIMM battery

A NVDIMM battery is responsible for maintaining power to the NVDIMM module.

PCIe

A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF A300

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.

- For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.

- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct node:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy node* is the HA partner of the impaired node.

Check encryption key support and status - AFF A300

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.

ONTAP version	Run this command
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	Follow these steps…
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the impaired controller - AFF A300

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller displays…	Then…
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster configuration

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired node.

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 3: Controller is in a two-node MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired node.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

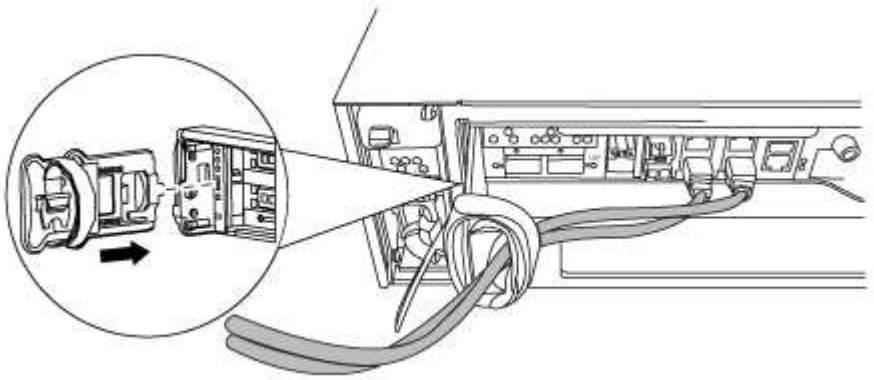

4. Loosen the thumbscrew on the cam handle on the controller module.

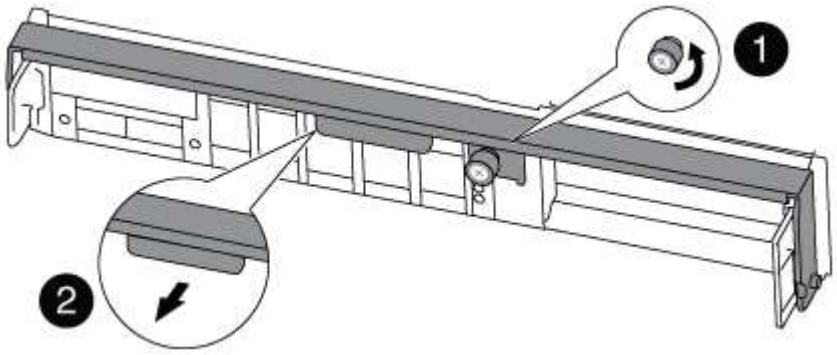

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 2: Replace the boot media - AFF A300

You must locate the boot media in the controller and follow the directions to replace it.

1. If you are not already grounded, properly ground yourself.
2. Locate the boot media using the following illustration or the FRU map on the controller module:

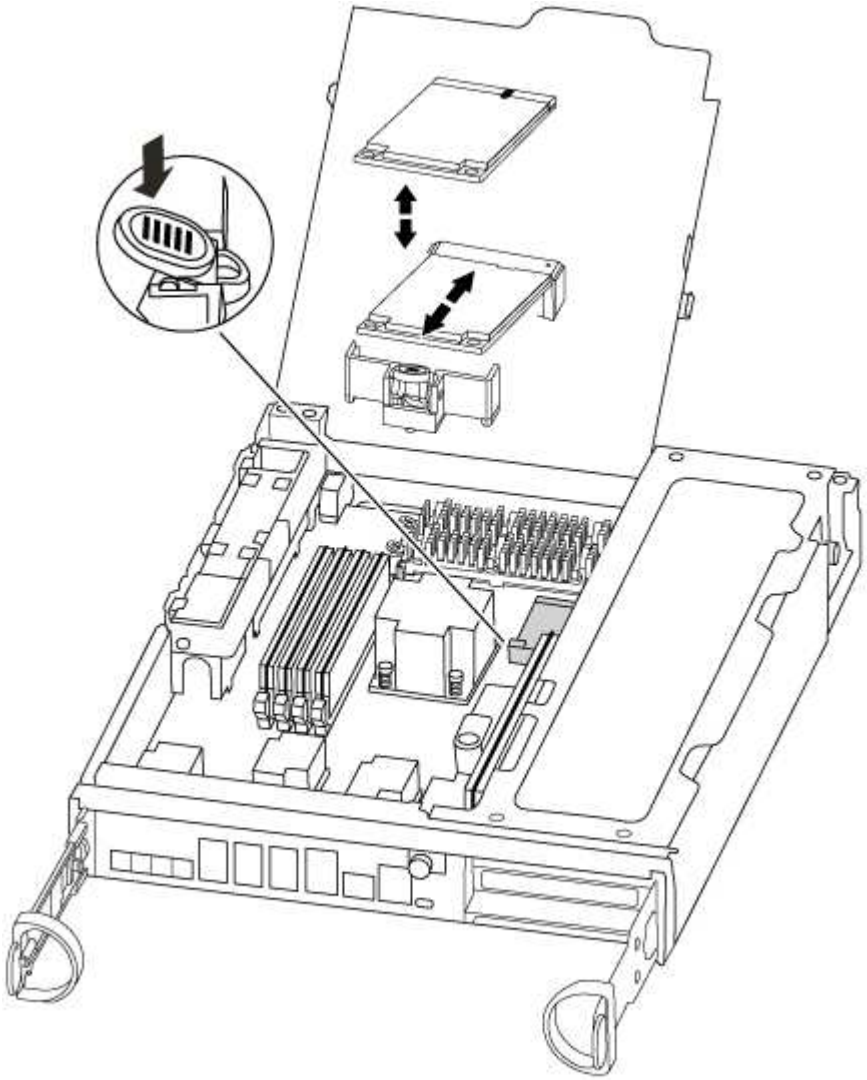

3. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

4. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
5. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

6. Push the boot media down to engage the locking button on the boot media housing.
7. Close the controller module cover.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

   The controller begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

6. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

   The system begins to boot and stops at the LOADER prompt.

7. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

The target port you configure is the target port you use to communicate with the impaired controller from the healthy controller during var file system restore with a network connection. You can also use the e0M port in this command.

- If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
  - filer_addr is the IP address of the storage system.
  - netmask is the network mask of the management network that is connected to the HA partner.
  - gateway is the gateway for the network.
  - dns_addr is the IP address of a name server on your network.
  - dns_domain is the Domain Name System (DNS) domain name.

    If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

8. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

The changes will be implemented when the system is booted.

Boot the recovery image - AFF A300

The procedure for booting the impaired controller from the recovery image depends on whether the system is in a two-controller MetroCluster configuration.

Option 1: Most systems

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

This procedure applies to systems that are not in a two-node MetroCluster configuration.

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:
   - If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed
   - If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.
6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.

8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Option 2: Controller is in a two-node MetroCluster

You must boot the ONTAP image from the USB drive and verify the environmental variables.

This procedure applies to systems in a two-node MetroCluster configuration.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. After the image is installed, start the restoration process:
   a. Press `n` when prompted to restore the backup configuration.
   b. Press `y` when prompted to reboot to start using the newly installed software.

      You should be prepared to interrupt the boot process when prompted.

4. As the system boots, press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message., and when the Boot Menu is displayed select option 6.
5. Verify that the environmental variables are set as expected.
   a. Take the node to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
   e. Reboot the node.

Switch back aggregates in a two-node MetroCluster configuration - AFF A300

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
**********************************************************************
**************
* Select option "(1) Normal Boot." to complete the recovery process.
*
**********************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF A300

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF A300

To replace the chassis, you must move the power supplies, fans, and controller modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - - AFF A300

To replace the chassis, you must shutdown the controllers.

Option 1: Shut down the controller

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

```
Warning: Are you sure you want to halt node <node_name>? {y|n}:
```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Controllers are in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace hardware - AFF A300

Move the power supplies, fans, and controller modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Move a power supply

Moving out a power supply when replacing a chassis involves turning off, disconnecting, and removing the power supply from the old chassis and installing and connecting it on the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Press down the release latch on the power supply cam handle, and then lower the cam handle to the fully open position to release the power supply from the mid plane.

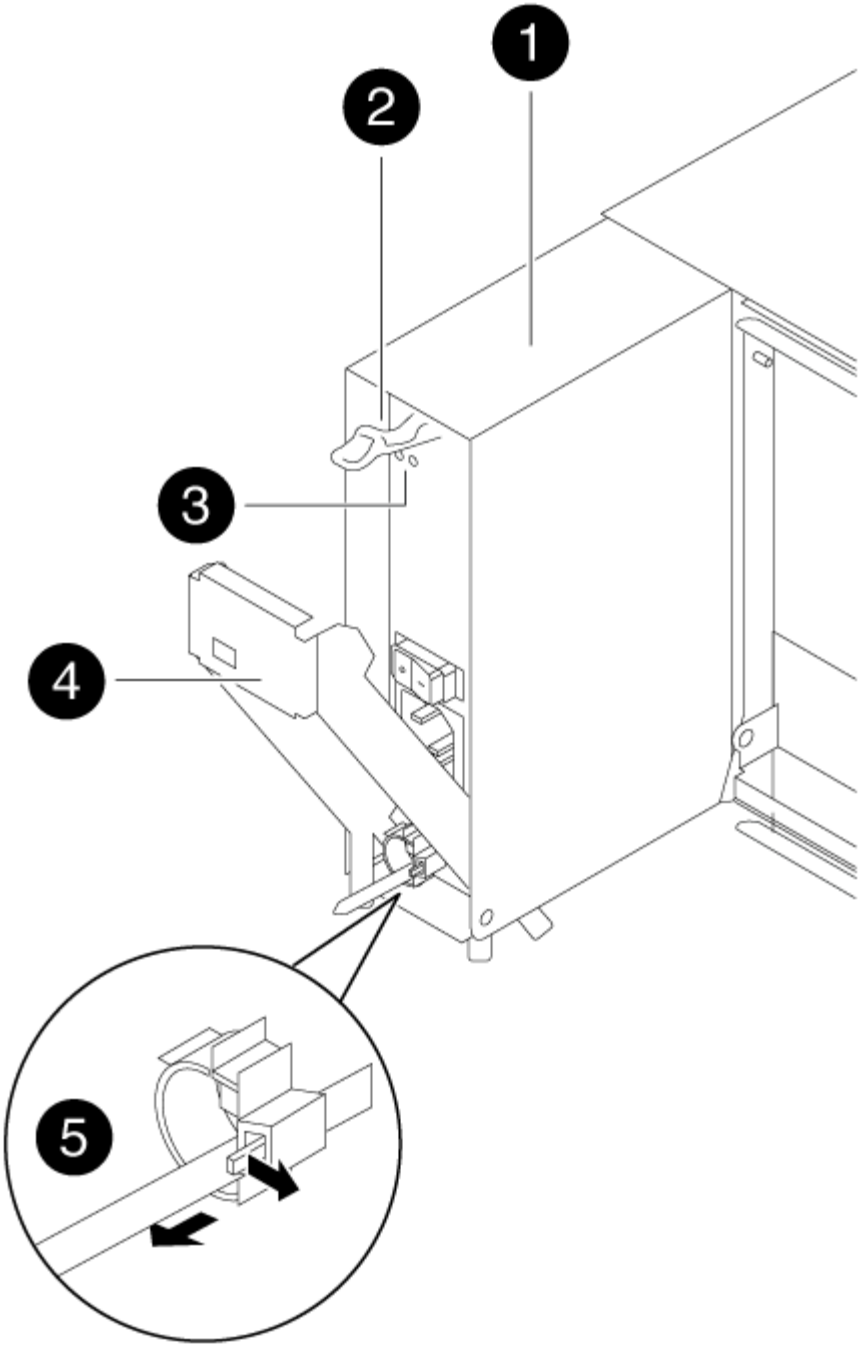

1	Power supply
2	Cam handle release latch
3	Power and Fault LEDs
4	Cam handle

5	Power cable locking mechanism

4. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Push firmly on the power supply cam handle to seat it all the way into the chassis, and then push the cam handle to the closed position, making sure that the cam handle release latch clicks into its locked position.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

   Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

Step 2: Move a fan

Moving out a fan module when replacing the chassis involves a specific sequence of tasks.

1. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
2. Press down the release latch on the fan module cam handle, and then pull the cam handle downward.

   The fan module moves a little bit away from the chassis.

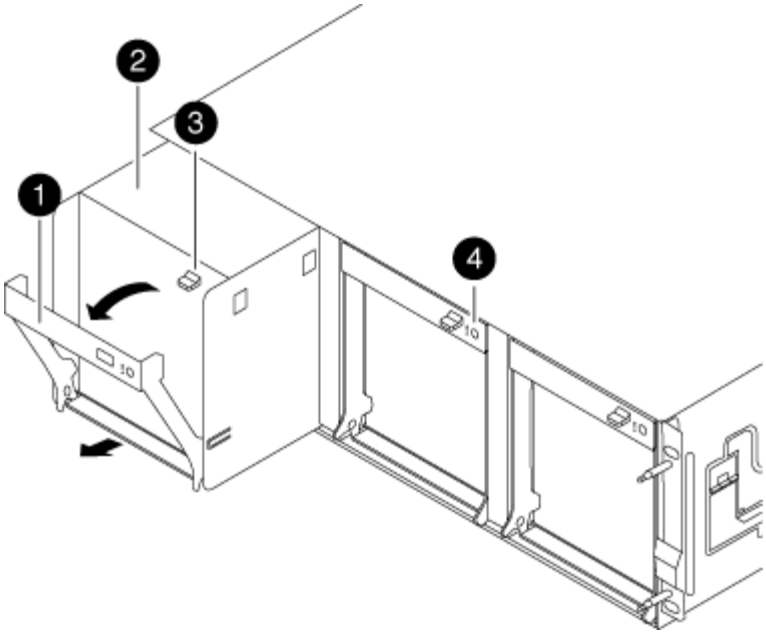

1	Cam handle
2	Fan module
3	Cam handle release latch
4	Fan module Attention LED

3. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

4. Set the fan module aside.
5. Repeat the preceding steps for any remaining fan modules.
6. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
7. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

8. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The fan LED should be green after the fan is seated and has spun up to operational speed.

9. Repeat these steps for the remaining fan modules.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Step 3: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

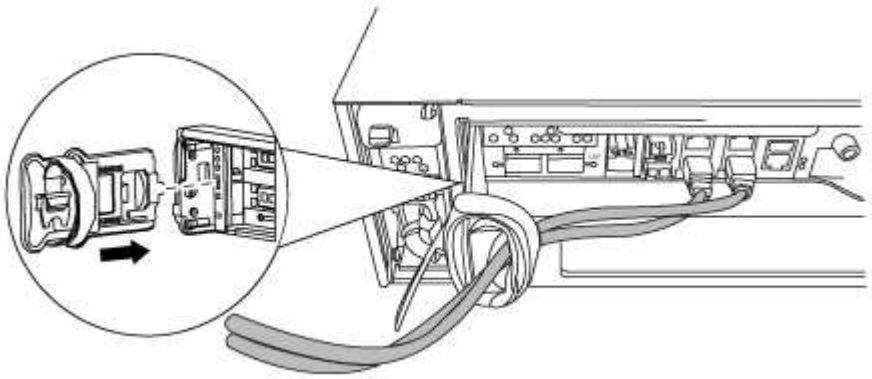

4. Loosen the thumbscrew on the cam handle on the controller module.

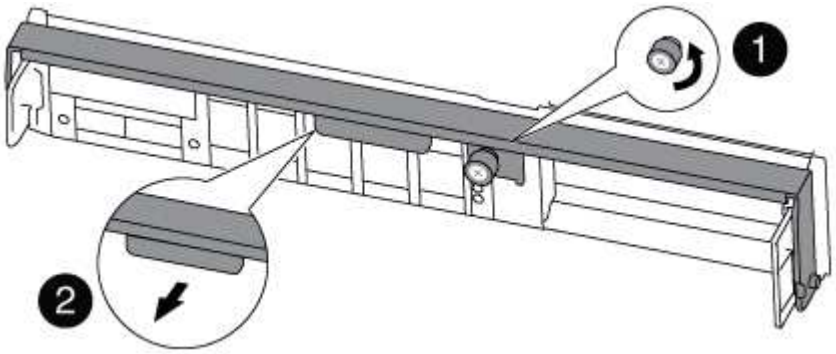

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

6. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 4: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.

   If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller

After you install the controller module and any other components into the new chassis, you must boot your system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. If you are not already grounded, properly ground yourself.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the console to the controller module, and then reconnect the management port.
4. Repeat the preceding steps if there is a second controller to install in the new chassis.
5. Complete the installation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Repeat the preceding steps for the second controller module in the new chassis.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reinstall the blanking panel and then go to the next step.

6. Connect the power supplies to different power sources, and then turn them on.
7. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

      If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

   b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - AFF A300

You must verify the HA state of the chassis, switch back aggregates, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis *HA-state*`

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mcc-2n`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. The next step depends on your system configuration.

If your system is in…	Then…
A stand-alone configuration	a. Exit Maintenance mode: `halt` b. Go to Step 3: Return the failed part to NetApp.
An HA pair with a second controller module	Exit Maintenance mode: `halt` The LOADER prompt appears.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller module

Overview of controller module replacement - AFF A300

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- This procedure includes steps for automatically or manually reassigning drives to the *replacement* controller, depending on your system’s configuration.

  You should perform the drive reassignment as directed in the procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- Any PCIe cards moved from the old controller module to the new controller module or added from existing customer site inventory must be supported by the replacement controller module.

  NetApp Hardware Universe

- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Open the controller module

To replace the controller module, you must first remove the old controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

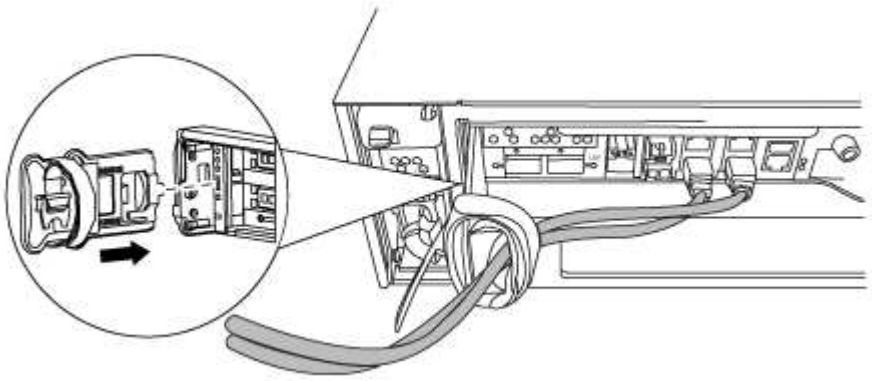

4. If you left the SFP modules in the system after removing the cables, move them to the new controller module.
5. Loosen the thumbscrew on the cam handle on the controller module.

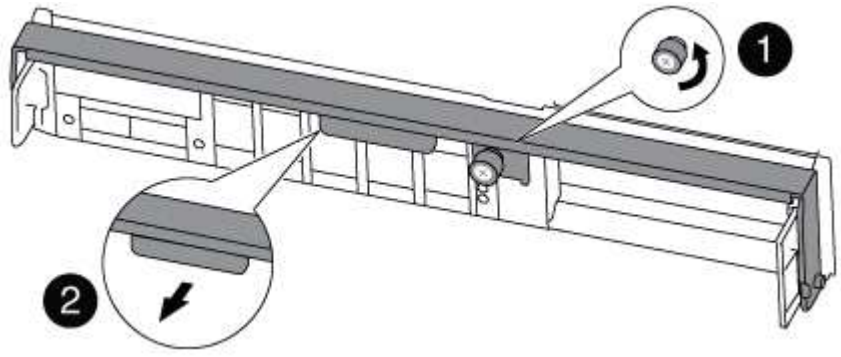

1	Thumbscrew
2	Cam handle

6. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 2: Move the boot device

You must locate the boot media and follow the directions to remove it from the old controller and insert it in the new controller.

1. Locate the boot media using the following illustration or the FRU map on the controller module:

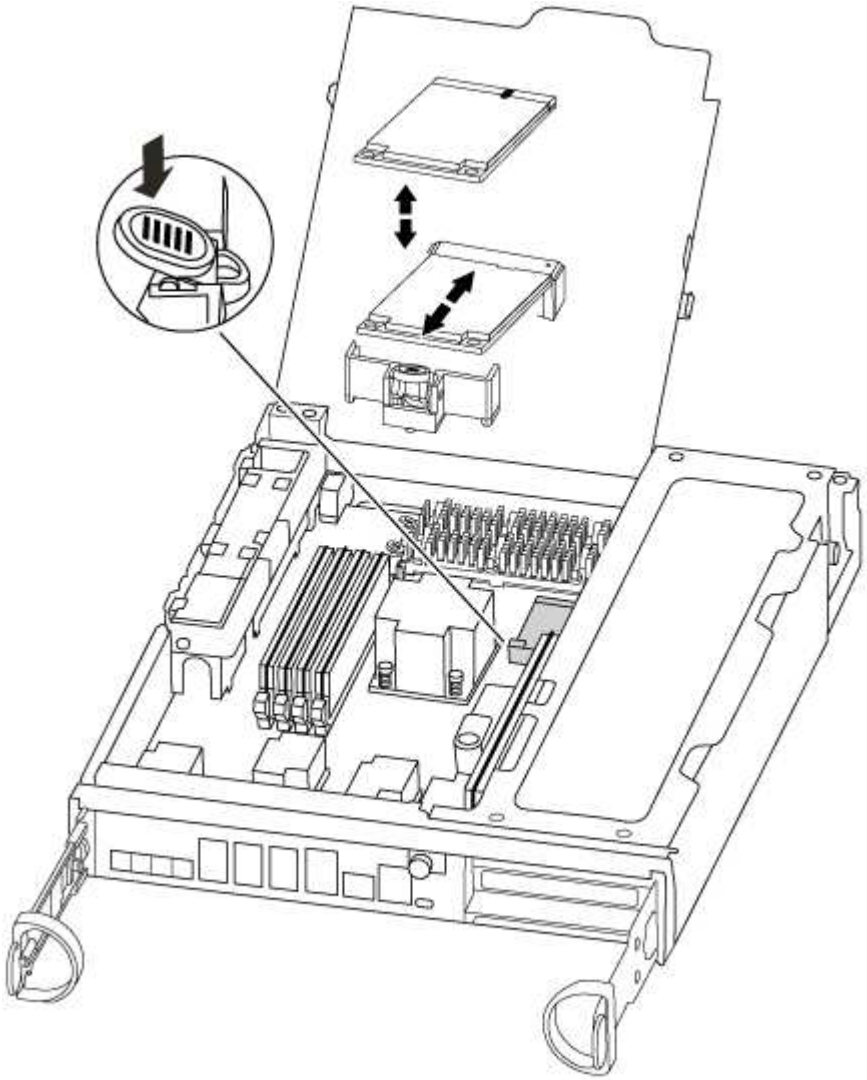

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the NVMEM battery

To move the NVMEM battery from the old controller module to the new controller module, you must perform a specific sequence of steps.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Open the CPU air duct and locate the NVMEM battery.

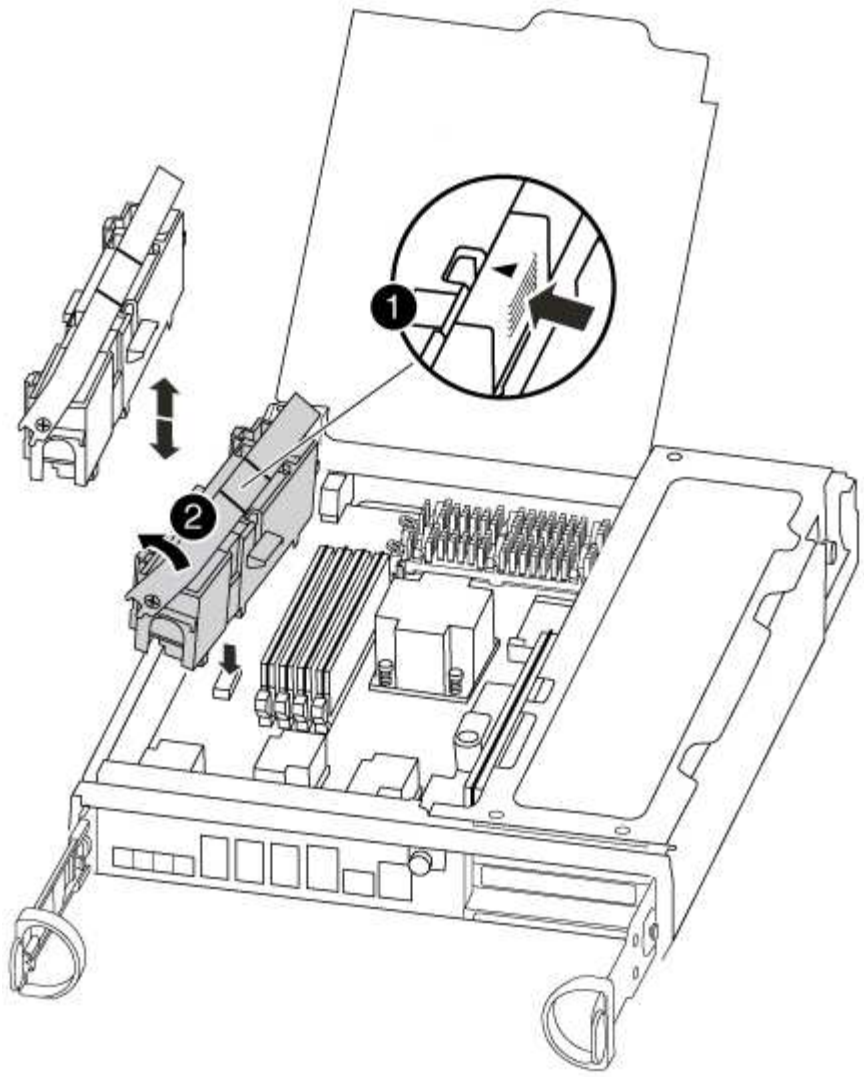

1	Battery lock tab
2	NVMEM battery pack

3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
4. Remove the battery from the controller module and set it aside.

Step 4: Move the DIMMs

To move the DIMMs, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

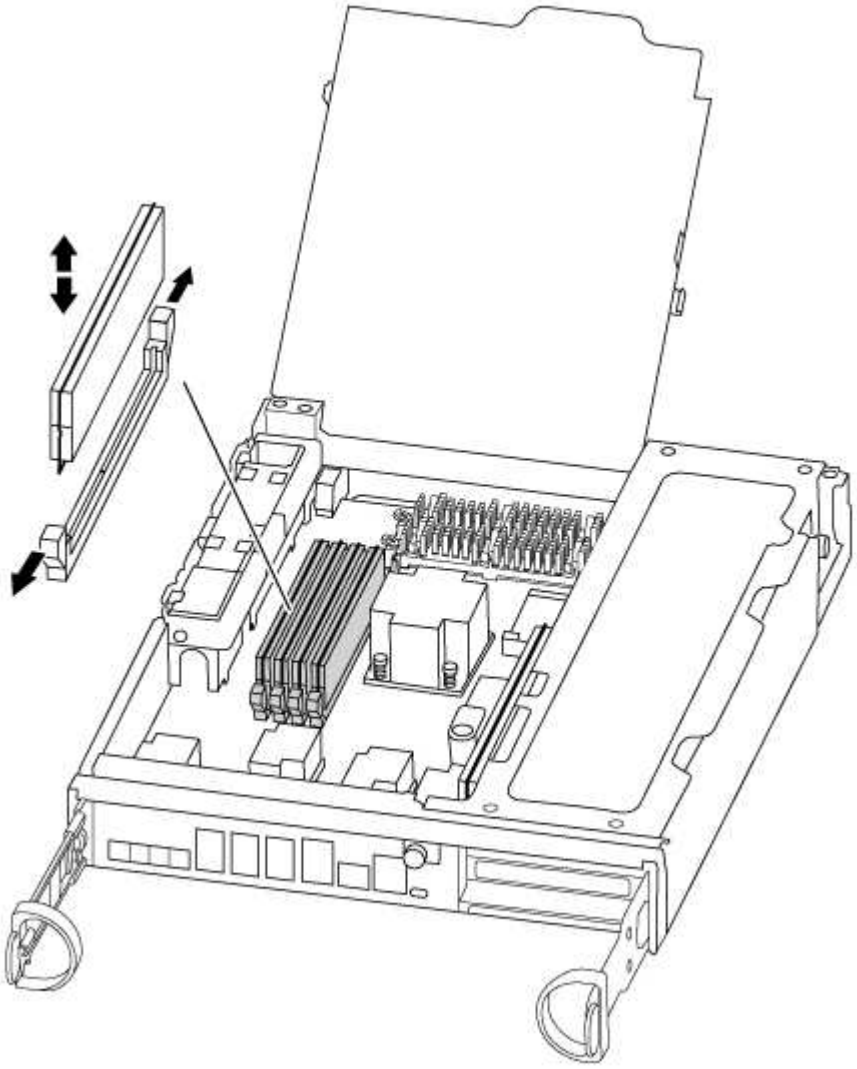

4. Locate the slot where you are installing the DIMM.
5. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Repeat these steps for the remaining DIMMs.
7. Move the NVMEM battery to the replacement controller module.
8. Align the tab or tabs on the battery holder with the notches in the controller module side, and then gently push down on the battery housing until the battery housing clicks into place.

Step 5: Move a PCIe card

To move PCIe cards, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

You must have the new controller module ready so that you can move the PCIe cards directly from the old controller module to the corresponding slots in the new one.

1. Loosen the thumbscrew on the controller module side panel.
2. Swing the side panel off the controller module.

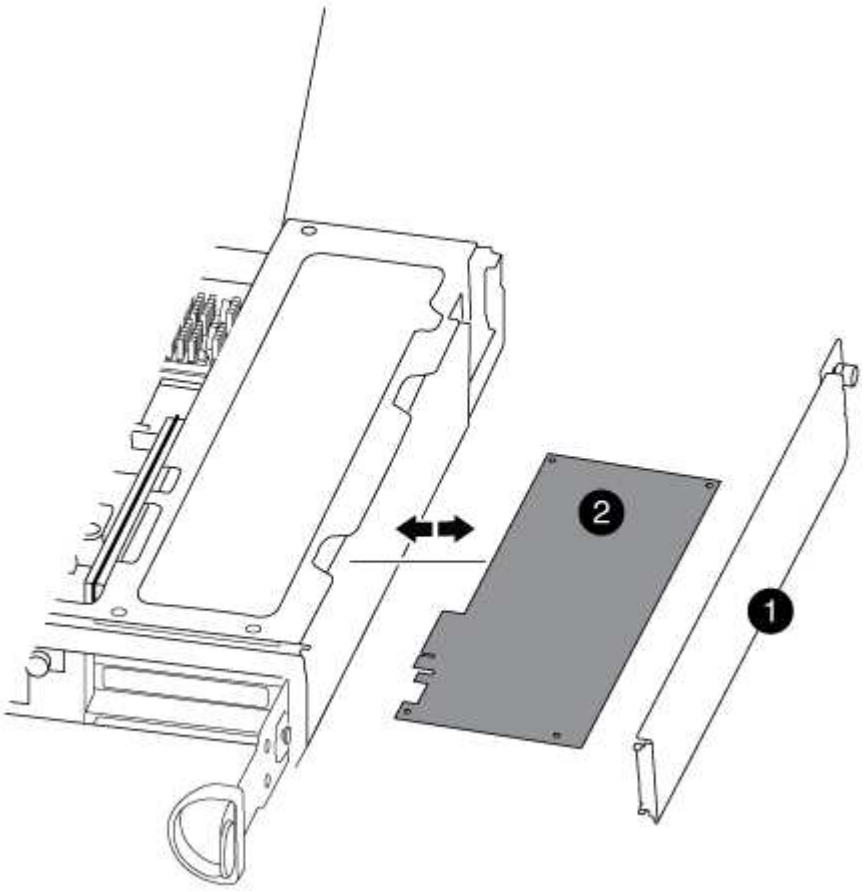

1	
Side panel	
2	
PCIe card	

3. Remove the PCIe card from the old controller module and set it aside.

   Make sure that you keep track of which slot the PCIe card was in.

4. Repeat the preceding step for the remaining PCIe cards in the old controller module.

5. Open the new controller module side panel, if necessary, slide off the PCIe card filler plate, as needed, and carefully install the PCIe card.

   Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The card must be fully and evenly seated in the slot.

6. Repeat the preceding step for the remaining PCIe cards that you set aside.
7. Close the side panel and tighten the thumbscrew.

Step 6: Install the controller

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

> The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, close the CPU air duct.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   > Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   > You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	`The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.` a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. When you see the message `Press Ctrl-C for Boot Menu`, press `Ctrl-C` to interrupt the boot process. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. Select the option to boot to Maintenance mode from the displayed menu.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process, and then press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. From the boot menu, select the option for Maintenance mode.

Important: During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
  You can safely respond `y` to these prompts.

Restore and verify the system configuration - AFF A300

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF A300

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must use the correct procedure for your configuration.

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* node and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* node is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* node, boot the node, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* node console and then, from the healthy node, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired node, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy node, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the node:
   a. From the healthy node, give back the replaced node's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* node takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.
7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* node should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Option 2: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `` `metrocluster node show -fields node-systemid`,dr-partner-systemid` ``

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

```
dr-group-id cluster         node                 node-systemid dr-
partner-systemid
 ----------- -------------------- ------------------- ------------
------------------
 1           Cluster_A            Node_A_1             536872914
118073209
 1           Cluster_B            Node_B_1             118073209
536872914
 2 entries were displayed.
```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

```
Local System ID: 118065481
    ...
    ...
```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                POOL   SERIAL NUMBER  HOME
-------    -------------        -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:

a. Change to the advanced privilege level: `set -privilege advanced`

   You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

b. Verify that the coredumps are saved: `system node run -node *local-node-name* partner savecore`

   If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node *local-node-name* partner savecore -s command`.</info>.

c. Return to the admin privilege level: `set -privilege admin`

7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
8. Boot the *replacement* node: `boot_ontap`
9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`
10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`

c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - AFF A300

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement node in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`

b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A300

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Open the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

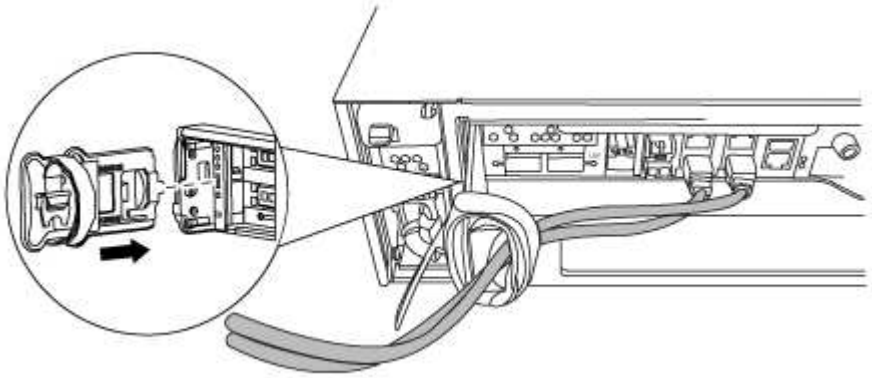

4. Loosen the thumbscrew on the cam handle on the controller module.

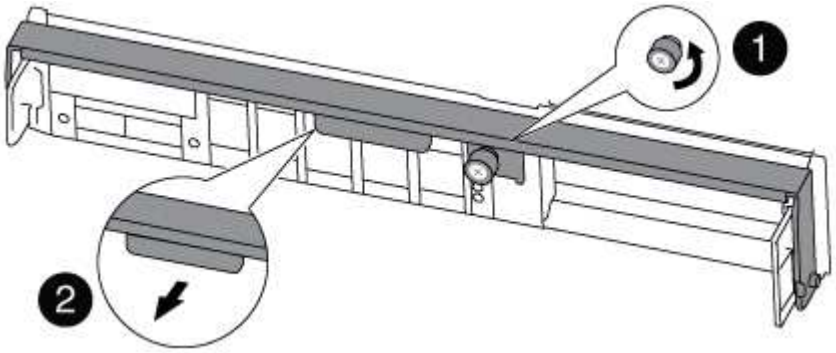

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED on the controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

   

3. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
4. Unplug the battery:

   The NVMEM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after Data ONTAP has successfully booted.

     a. Open the CPU air duct and locate the NVMEM battery.

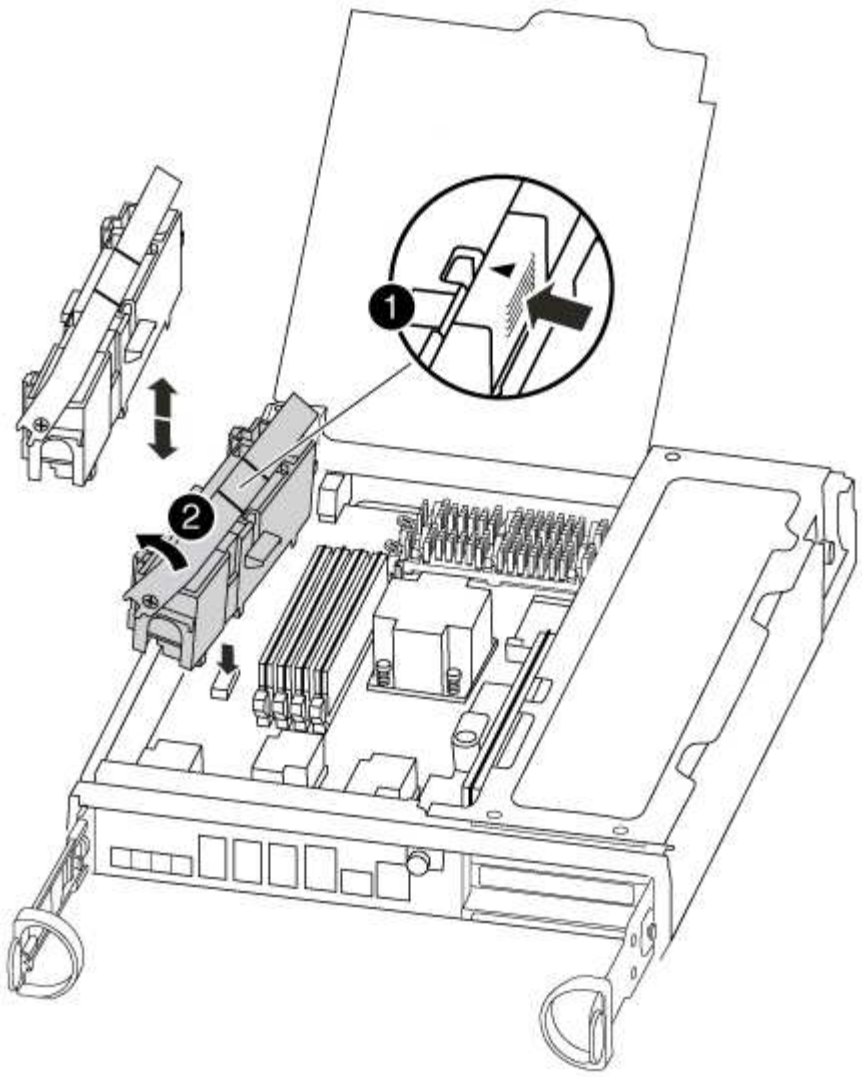

1	NVMEM battery lock tab
2	NVMEM battery

b. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.

c. Wait a few seconds, and then plug the battery back into the socket.

5. Return to step 2 of this procedure to recheck the NVMEM LED.
6. Locate the DIMMs on your controller module.
7. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
8. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

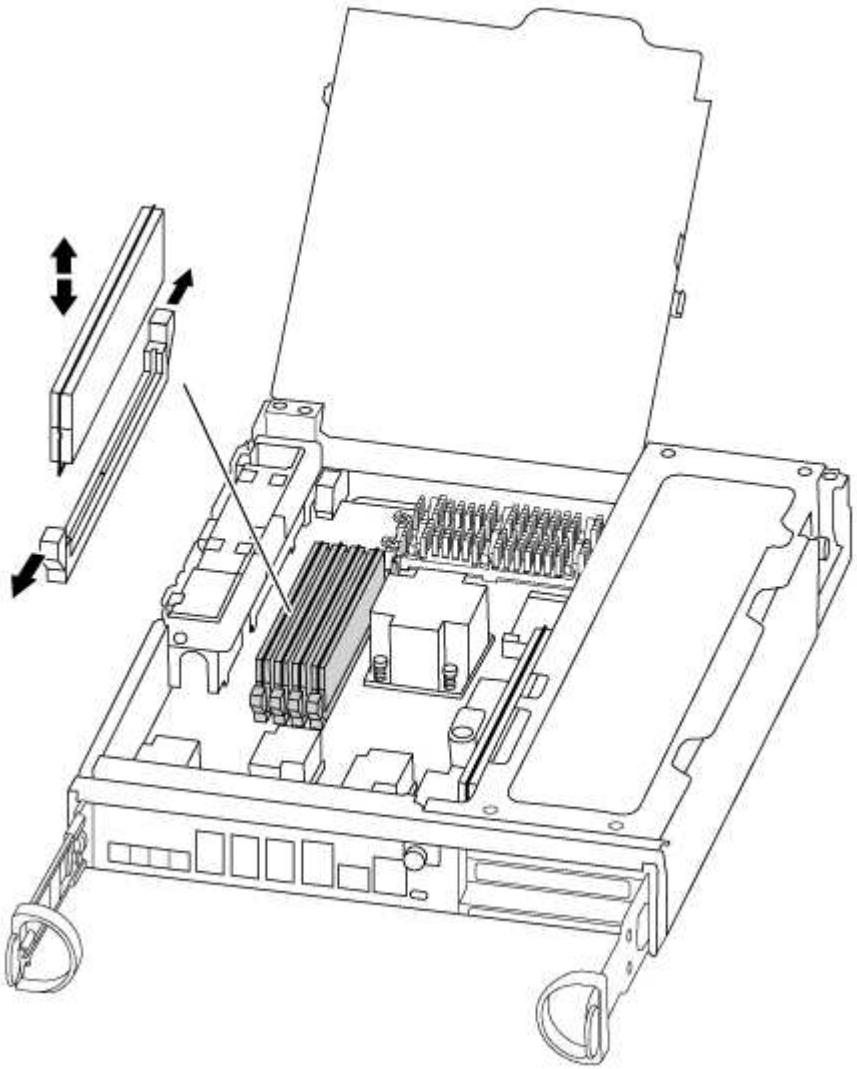

9. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

10. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

    The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

    Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

11. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
12. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

    Make sure that the plug locks down onto the controller module.

13. Close the controller module cover.

Step 4: Reinstall the controller

After you replace a component within the controller module, you must reinstall the controller module in the system chassis.

1. If you are not already grounded, properly ground yourself.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:

   The controller module begins to boot as soon as it is fully seated in the chassis.

   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Tighten the thumbscrew on the cam handle on back of the controller module.
   c. If you have not already done so, reinstall the cable management device.
   d. Bind the cables to the cable management device with the hook and loop strap.

Step 5: (Two-node MetroCluster only): Switch back aggregates

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a fan - AFF A300

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press down the release latch on the fan module cam handle, and then pull the cam handle downward.

   The fan module moves a little bit away from the chassis.

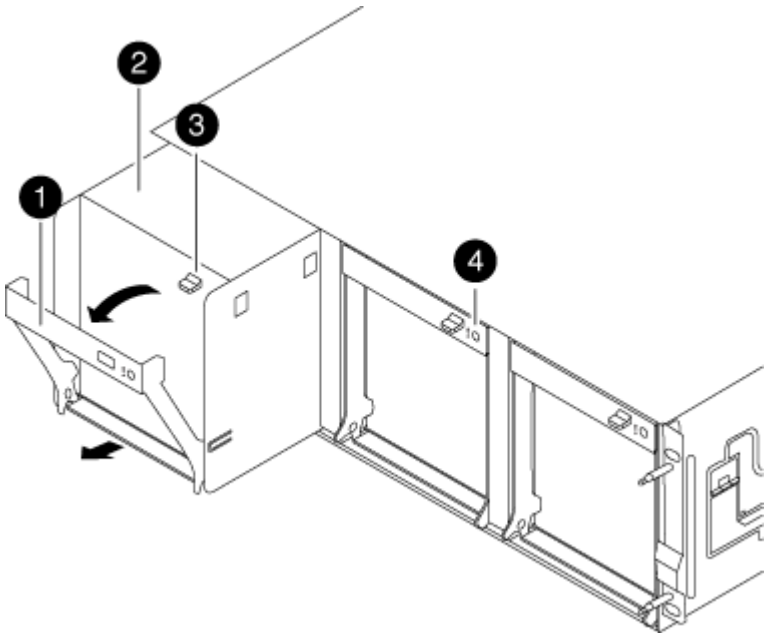

1	Cam handle
2	Fan module
2	Cam handle release latch

4	Fan module Attention LED

5. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

6. Set the fan module aside.
7. Insert the replacement fan module into the chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The fan LED should be green after the fan is seated and has spun up to operational speed.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVMEM battery - AFF A300

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Open the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

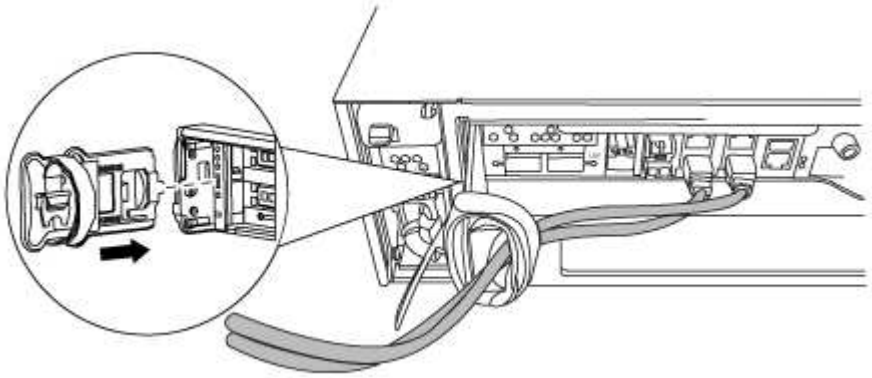

4. Loosen the thumbscrew on the cam handle on the controller module.

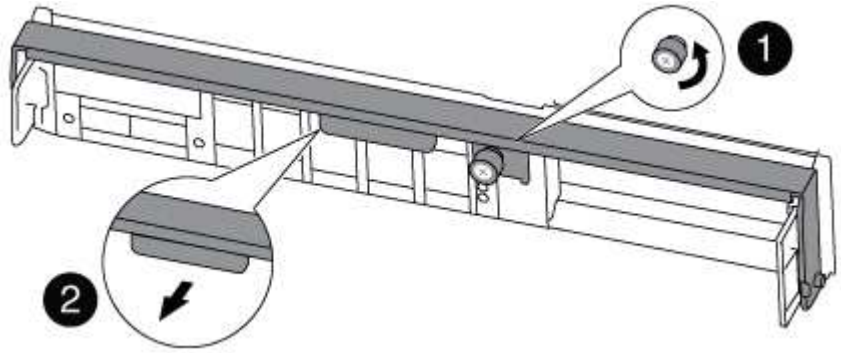

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery in your system, you must remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

> The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

3. Open the CPU air duct and locate the NVMEM battery.

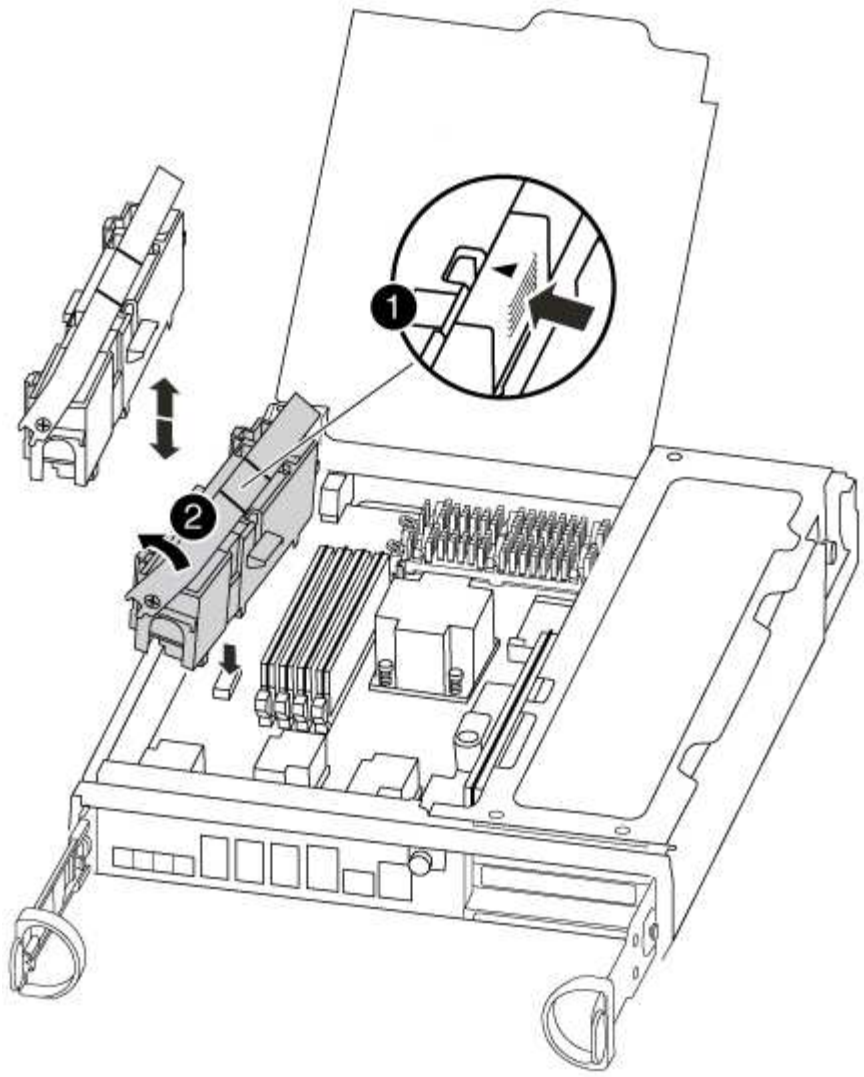

1	Battery lock tab
2	NVMEM battery pack

4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Remove the replacement battery from its package.
6. Align the tab or tabs on the battery holder with the notches in the controller module side, and then gently push down on the battery housing until the battery housing clicks into place.
7. Close the CPU air duct.

   Make sure that the plug locks down to the socket.

Step 4: Reinstall the controller

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. If you are not already grounded, properly ground yourself.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:

   The controller module begins to boot as soon as it is fully seated in the chassis.

   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Tighten the thumbscrew on the cam handle on back of the controller module.
   c. If you have not already done so, reinstall the cable management device.
   d. Bind the cables to the cable management device with the hook and loop strap.

Step 5: (two-node MetroCluster only): Switch back aggregates

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe card - AFF A300

To replace a PCIe card, you must perform a specific sequence of tasks.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Open the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

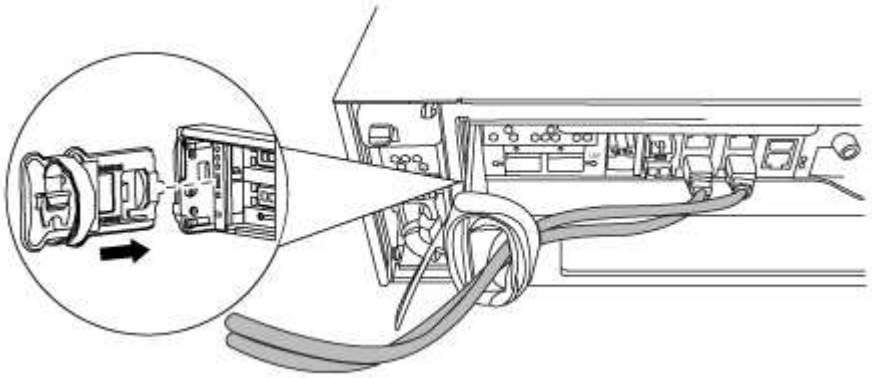

4. Loosen the thumbscrew on the cam handle on the controller module.

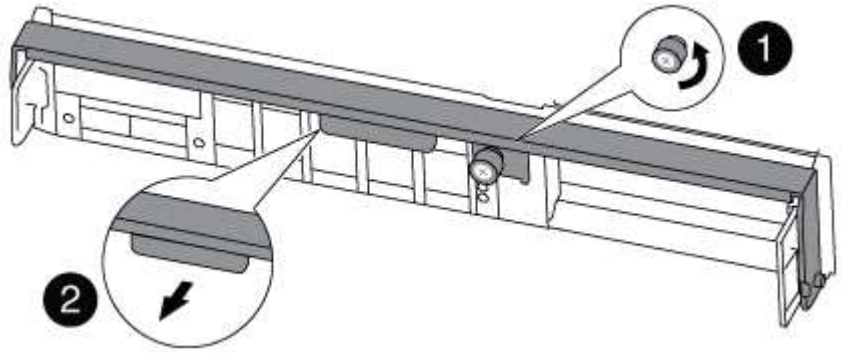

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace a PCIe card

To replace a PCIe card, locate it within the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Loosen the thumbscrew on the controller module side panel.
3. Swing the side panel off the controller module.

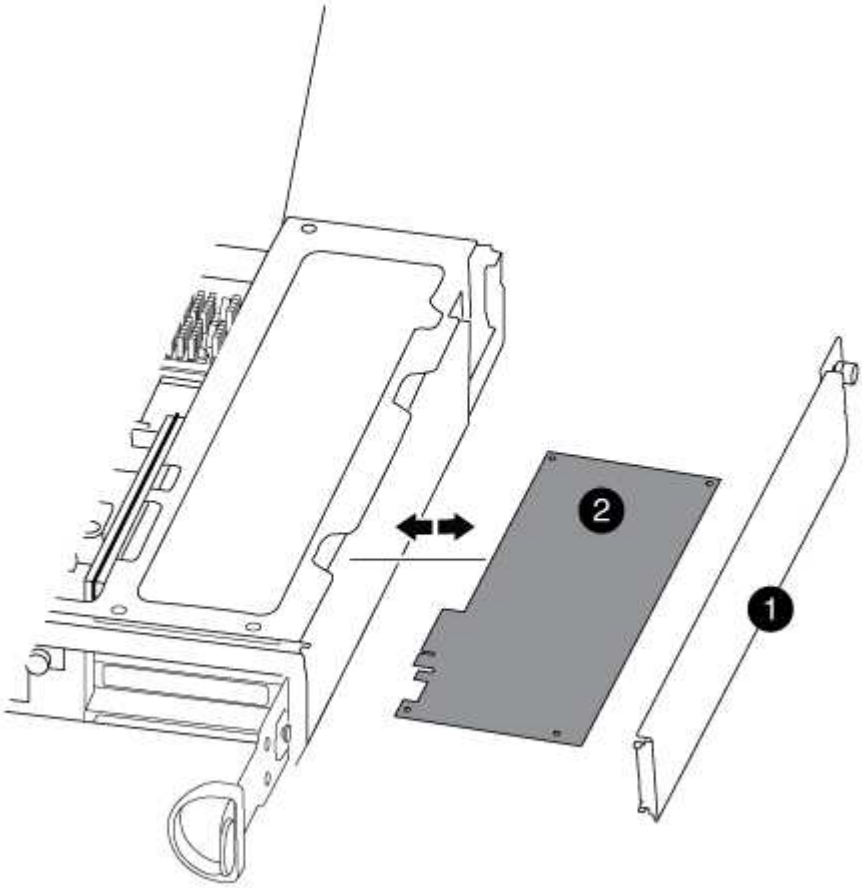

1	Side panel
2	PCIe card

4. Remove the PCIe card from the controller module and set it aside.
5. Install the replacement PCIe card.

   Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The PCIe card must be fully and evenly seated in the slot.

   If you are installing a card in the bottom slot and cannot see the card socket well, remove the top card so that you can see the card socket, install the card, and then reinstall the card you removed from the top slot.

6. Close the side panel and tighten the thumbscrew.

Step 4: Reinstall the controller

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. If you are not already grounded, properly ground yourself.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:

   The controller module begins to boot as soon as it is fully seated in the chassis.

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. If you have not already done so, reconnect the cables to the controller module. d. Bind the cables to the cable management device with the hook and loop strap.

If your system is in…	Then perform these steps…
A two-node MetroCluster configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. If you have not already done so, reconnect the cables to the controller module. d. Bind the cables to the cable management device with the hook and loop strap. e. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.

5. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs or onboard ports, convert these ports to 10 GbE connections by using the nicadmin convert command from Maintenance mode.

   Be sure to exit Maintenance mode after completing the conversion.

6. Return the controller to normal operation:

If your system is in…	Issue this command from the partner's console…
An HA pair	`storage failover giveback -ofnode` *`impaired_node_name`*
A two-node MetroCluster configuration	Proceed to the next step. The MetroCluster switchback procedure is done in the next task in the replacement process.

7. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5 (two-node MetroCluster only): Switch back aggregate

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - AFF A300

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- The number of power supplies in the system depends on the model.
- Power supplies are auto-ranging.

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Press down the release latch on the power supply cam handle, and then lower the cam handle to the fully open position to release the power supply from the mid plane.

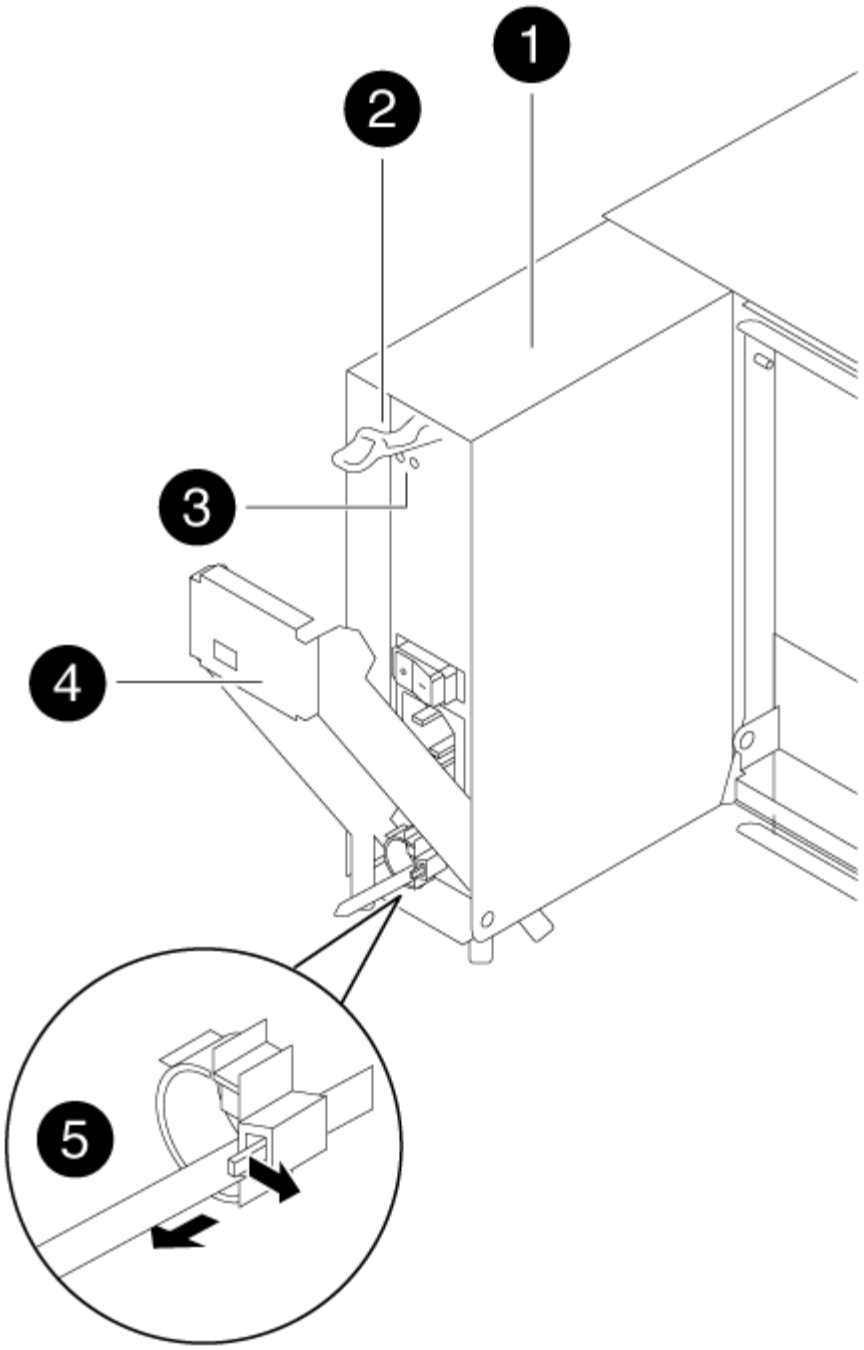

1	Power supply
2	Cam handle release latch
2	Power and Fault LEDs
4	Cam handle

5	Power cable locking mechanism

5. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.
7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Push firmly on the power supply cam handle to seat it all the way into the chassis, and then push the cam handle to the closed position, making sure that the cam handle release latch clicks into its locked position.
9. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

Once power is restored to the power supply, the status LED should be green.

1. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

   The power supply LEDs are lit when the power supply comes online.

2. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A300

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Open the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

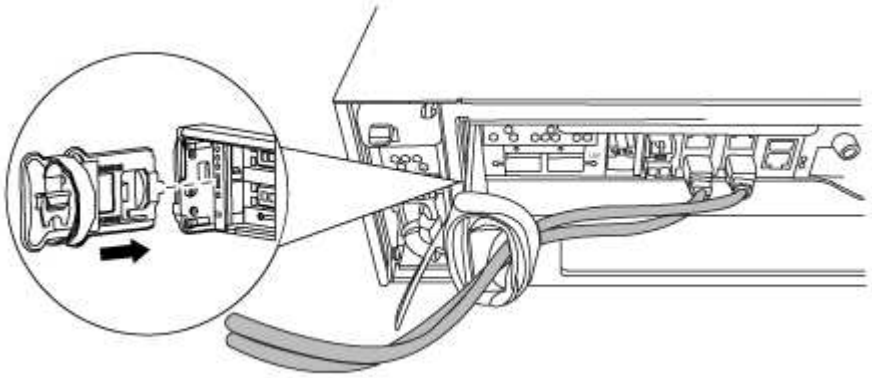

4. Loosen the thumbscrew on the cam handle on the controller module.

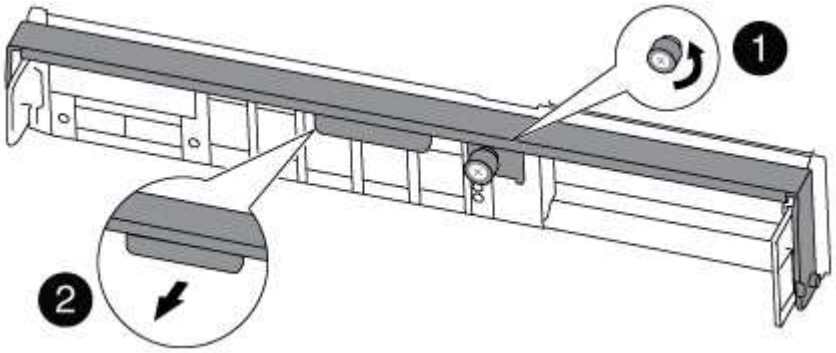

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the RTC Battery

To replace the RTC battery, locate them inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

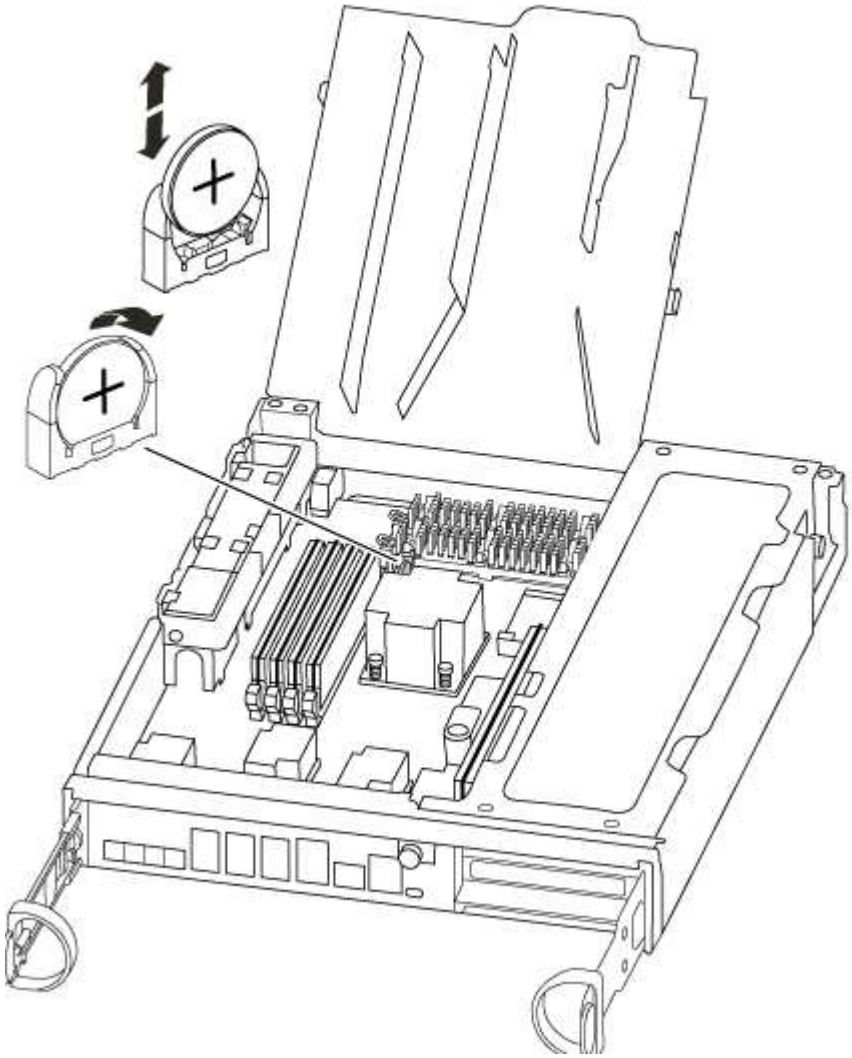

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Tighten the thumbscrew on the cam handle on back of the controller module.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF A320 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick guide - AFF A320

This guide gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Access the *Installation and Setup Instructions* PDF poster:

AFF A320 Installation and Setup Instructions

Video steps - AFF A320

The following video shows how to install and cable your new system.

[] | *https://img.youtube.com/vi/rQ-za_bli0Y?/maxresdefault.jpg*

Detailed guide - AFF A320

This guide gives detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

Prepare for installation

To install your AFF A320 system, you need to create an account, register the system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser
- A laptop or console with an RJ-45 connection and access to a Web browser

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register your system.

      NetApp Product Registration

4. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
100 GbE cable (QSF(28)	X66211A-05 (112-00595), 0.5m X66211A-1 (112-00573), 1m X66211A-2 (112-00574), 2m X66211A-5 (112-00574), 5m		Storage, cluster interconnect/HA, and Ethernet data (order-dependent)

Type of cable…	Part number and length	Connector type	For…
40 GbE cable	X66211A-1 (112-00573), 1m; X66211A-3 (112-00543),3m; X66211A-5 (112-00576), 5m		Storage, cluster interconnect/HA, and Ethernet data (order-dependent)
Ethernet cable - MPO	X66200-2 (112-00326), 2m X66250-5 (112-00328), 5m X66250-30 (112-00331), 30m		Ethernet cable (order dependent)
Optical cables	SR: X6553-R6 (112-00188), 2m X6554-R6 (112-00189), 15m X6537-R6 (112-00091), 30m LR: X66250-3 (112-00342), 2m X66260-5 (112-00344), 5m X66260-30 (112-00354), 30m		FC configurations (order-dependent)
RJ-45 (order dependent)	X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network
Micro-USB console cable	Not applicable		Console connection used during software setup if laptop or console does not support network discovery.
Power cables	Not applicable		Powering up the system

5. Download and complete the *Cluster configuration worksheet*.

   Cluster Configuration Worksheet

Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

1. Install the rail kits, as needed.

2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

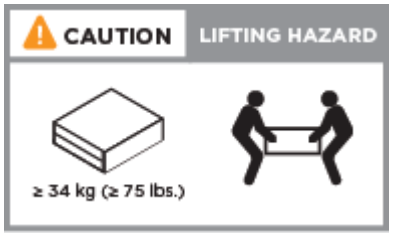

3. Attach cable management devices (as shown).

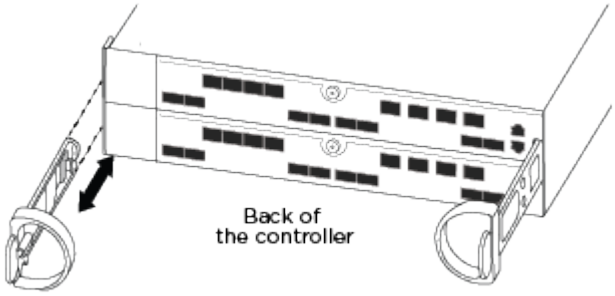

4. Place the bezel on the front of the system.

Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Cable a two-node switchless cluster

The optional data ports, optional NIC cards, and management ports on the controller modules are connected to switches. The cluster interconnect/HA ports are cabled on both controller modules.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

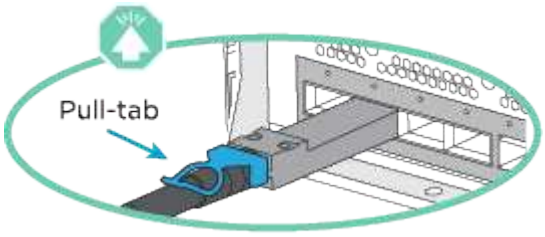

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

1. You can used the illustration or the step-by step instructions to complete the cabling between the controllers and to the switches:

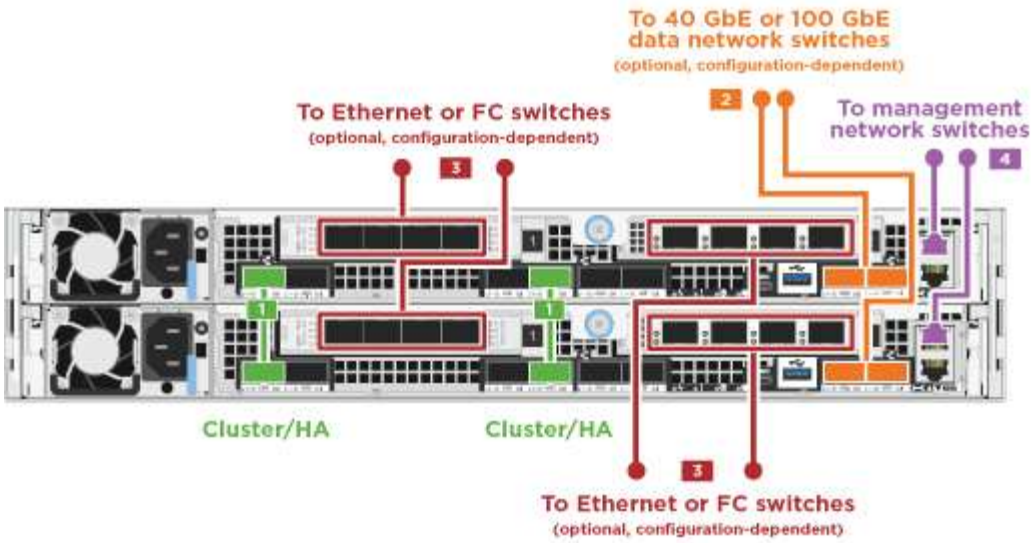

Step	Perform on each controller module
1	Cable the cluster/HA ports to each other with the 100 GbE (QSFP28) cable: • e0a to e0a • e0d to e0d 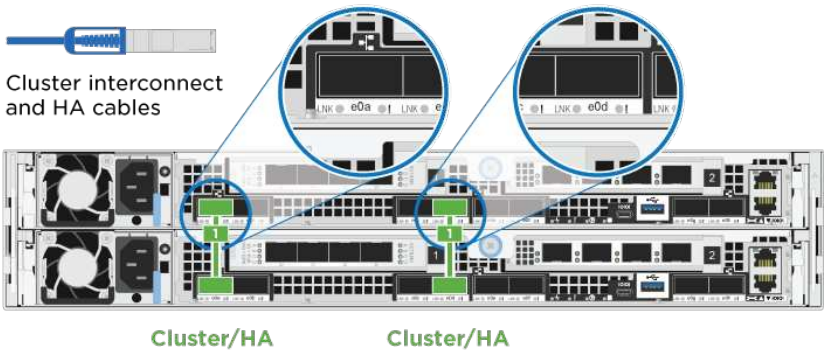

Step	Perform on each controller module
2	If you are using your onboard ports for a data network connection, connect the 100GbE or 40Gbe cables to the appropriate data network switches: • e0g and e0h 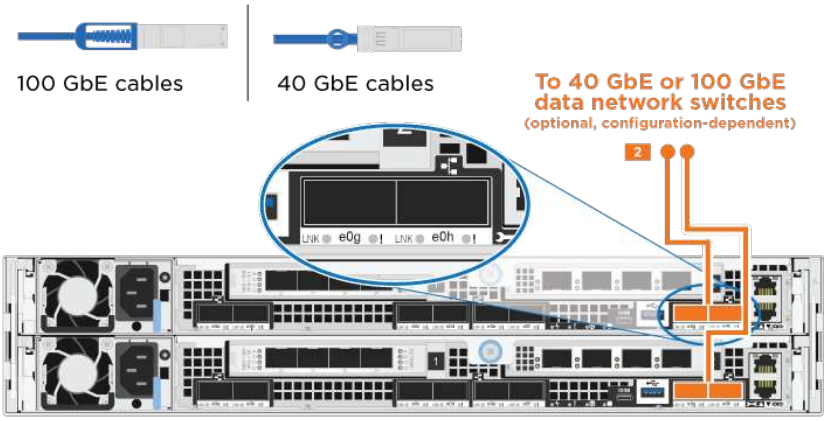
3	If you are using your NIC cards for Ethernet or FC connections, connect the NIC card(s) to the appropriate switches: 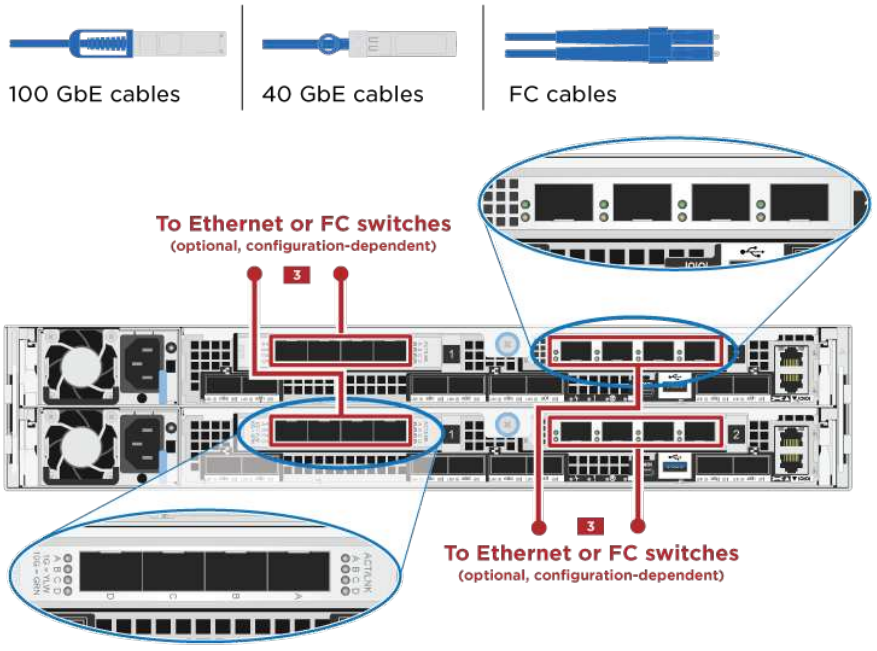

Step	Perform on each controller module
4	Cable the e0M ports to the management network switches with the RJ45 cables. 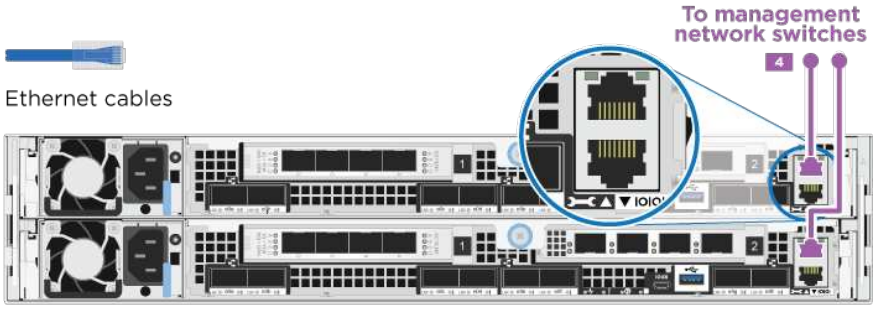
!	DO NOT plug in the power cords at this point.

2. Cable your storage: Cable controllers to drive shelves

Option 2: Cabling a switched cluster

The optional data ports, optional NIC cards, and management ports on the controller modules are connected to switches. The cluster interconnect/HA ports are cabled on to the cluster/HA switch.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

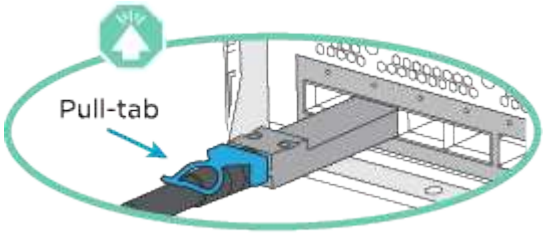

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

1. You can used the illustration or the step-by step instructions to complete the cabling between the controllers and to the switches:

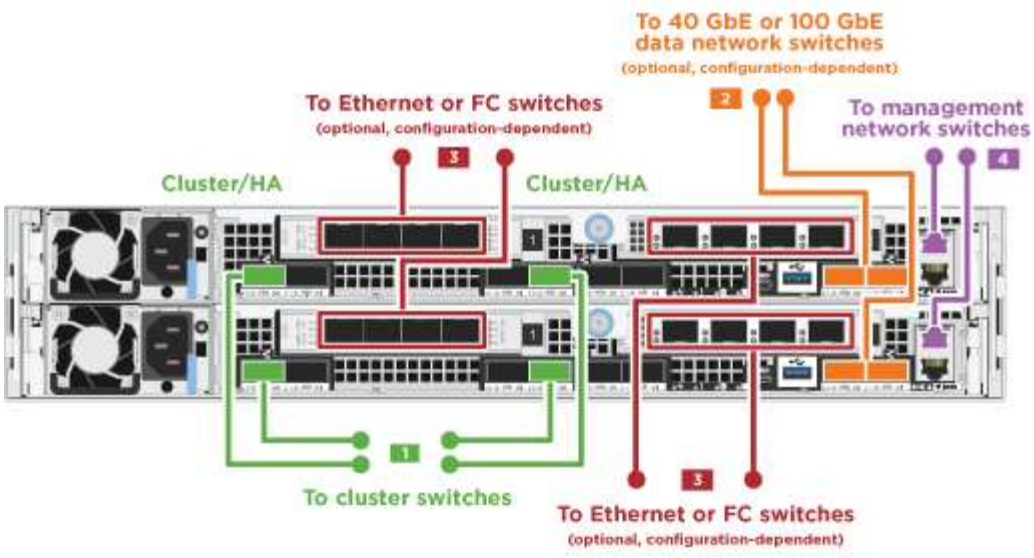

Step	Perform on each controller module
1	Cable the cluster/HA ports to the cluster/HA switch with the 100 GbE (QSFP28) cable: • e0a on both controllers to the cluster/HA switch • e0d on both controllers to the cluster/HA switch 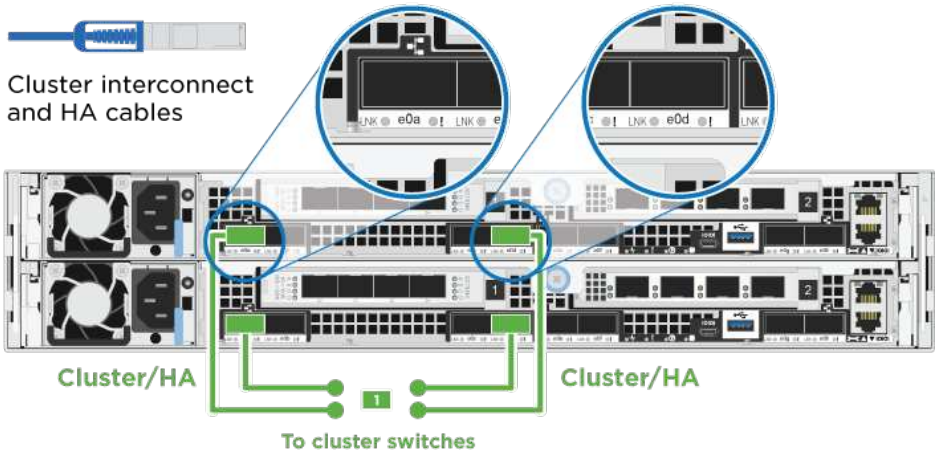

Step	Perform on each controller module
2	If you are using your onboard ports for a data network connection, connect the 100GbE or 40Gbe cables to the appropriate data network switches: • e0g and e0h 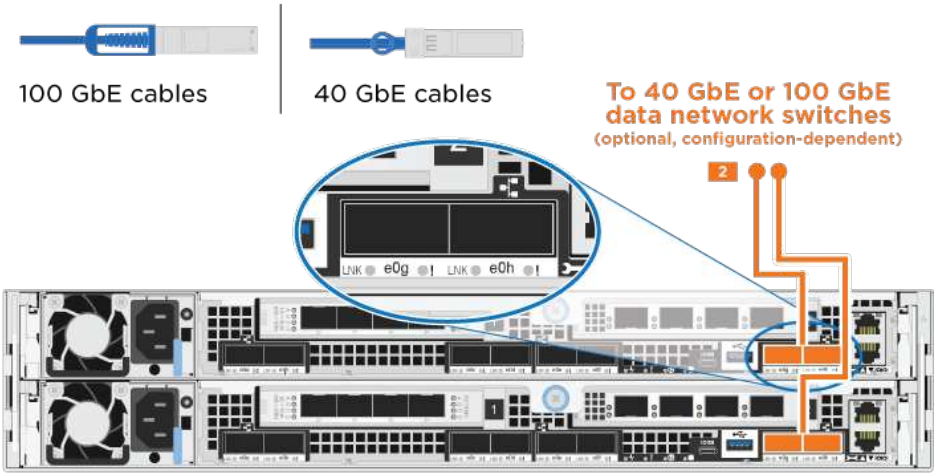
3	If you are using your NIC cards for Ethernet or FC connections, connect the NIC card(s) to the appropriate switches: 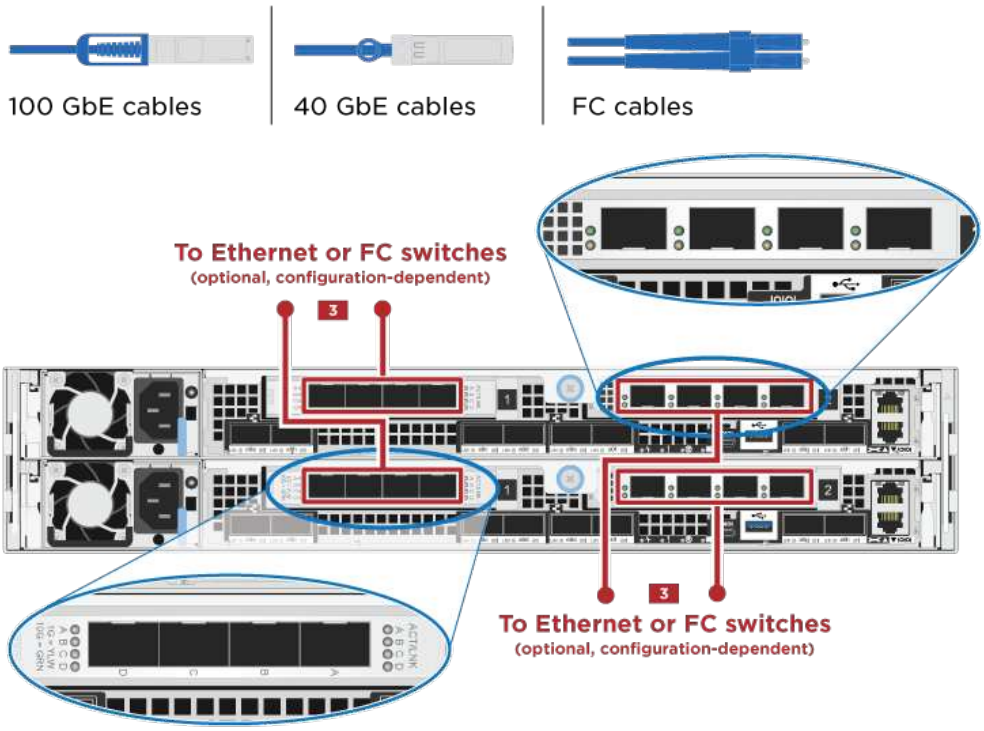

Step	Perform on each controller module
4	Cable the e0M ports to the management network switches with the RJ45 cables. 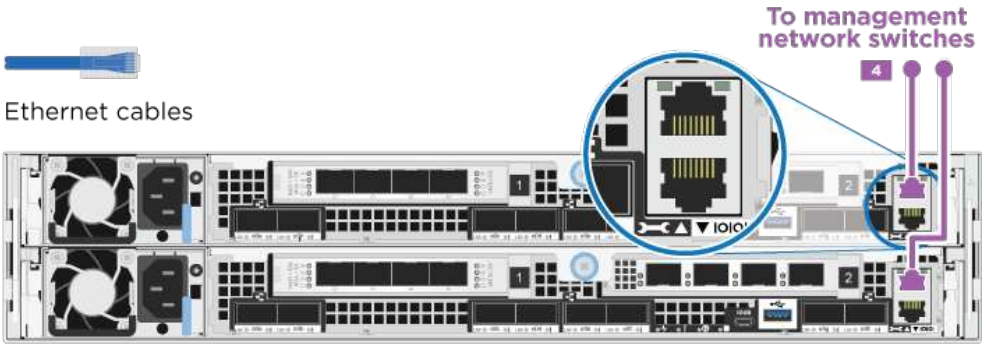
!	DO NOT plug in the power cords at this point.

2. Cable your storage: Cable controllers to drive shelves

Cable controllers to drive shelves

You must cable the controllers to your shelves using the onboard storage ports.

Option 1: Cable the controllers to a single drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

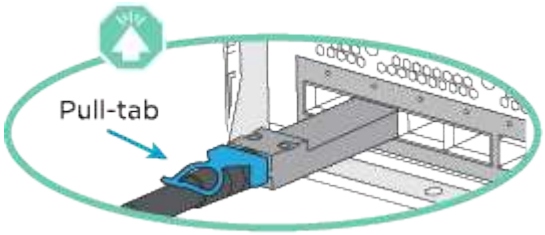

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

1. You can use the illustration or the step-by-step instructions to cable your controllers to a single shelf.

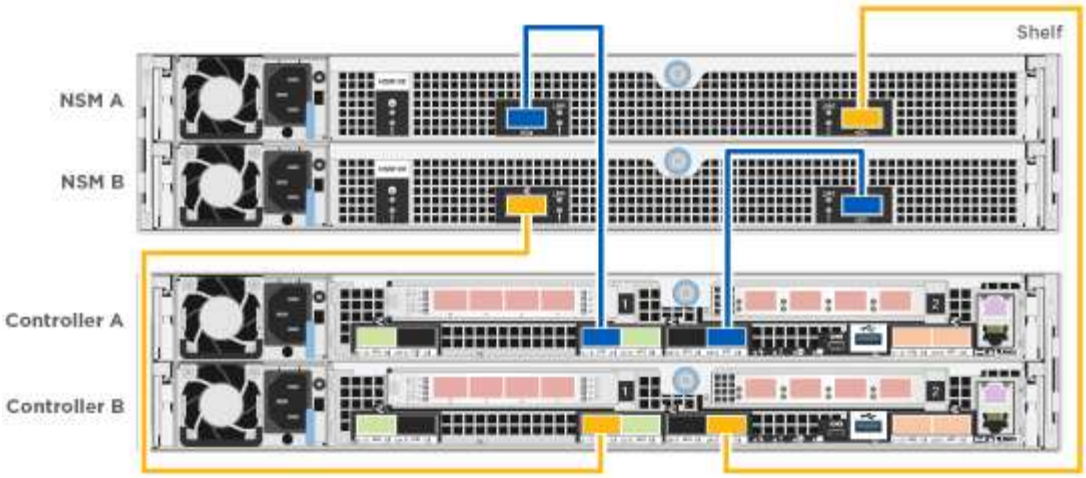

Step	Perform on each controller module
1	Cable controller A to the shelf 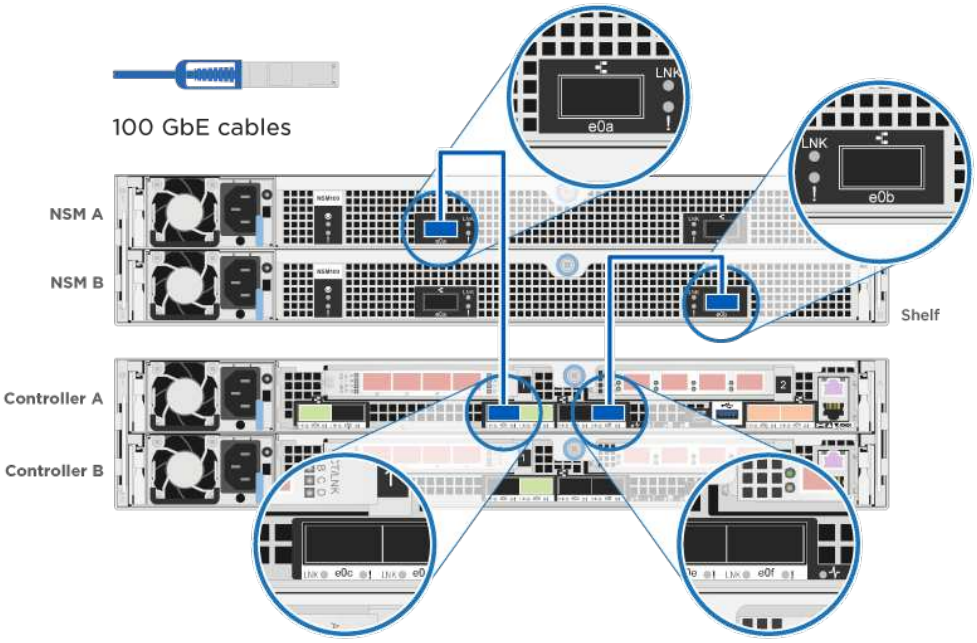

Step	Perform on each controller module
2	Cable controller B to the shelf: 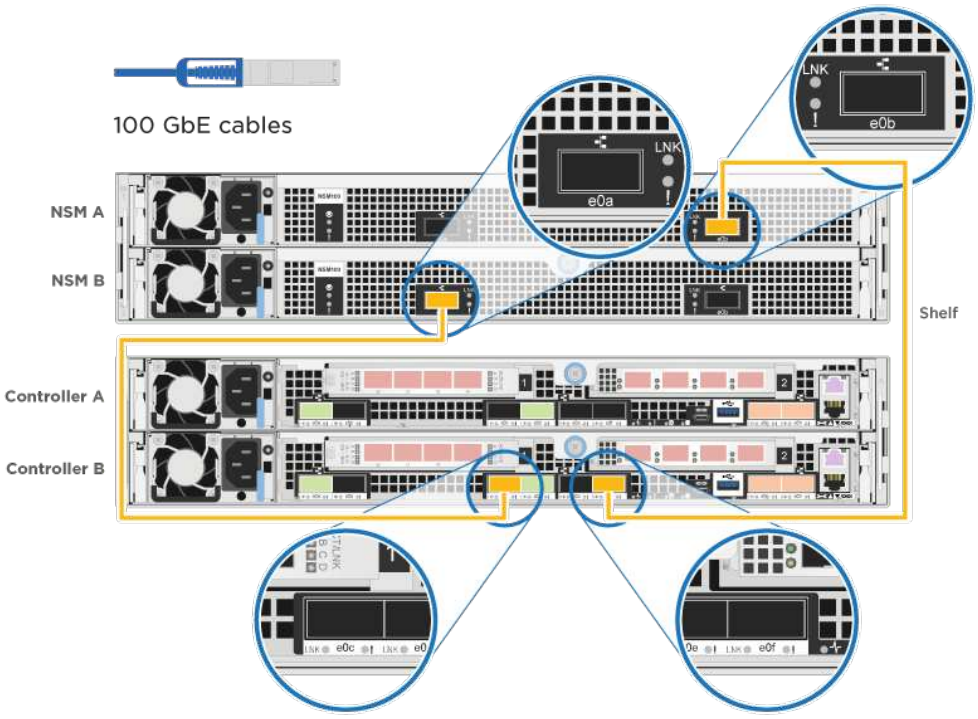

2. To complete setting up your system, see Complete system setup and configuration

Option 2: Cable the controllers to two drive shelves

You must cable each controller to the NSM modules on both NS224 drive shelves.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

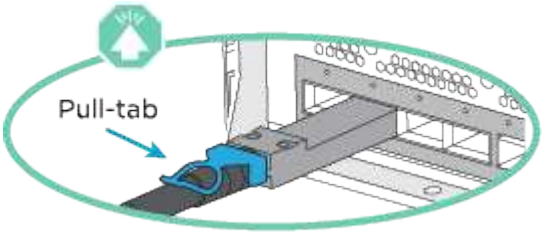

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

1. You can use the following illustration or the written steps to cable your controllers to two drive shelves.

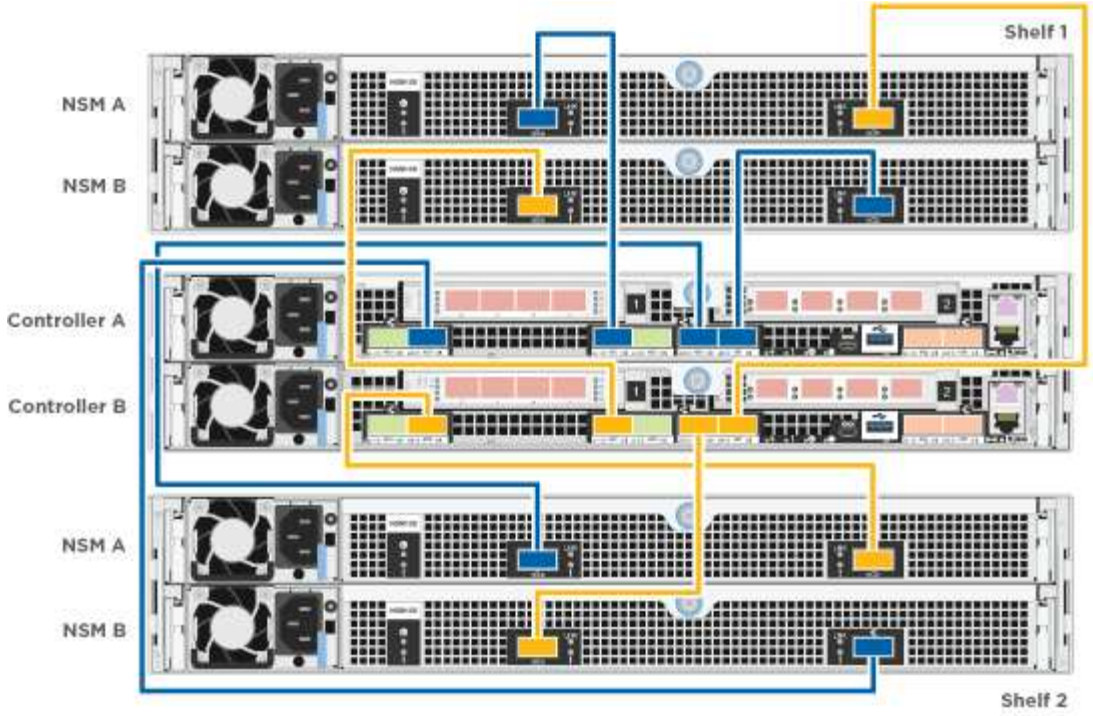

Step	Perform on each controller module
1	Cable controller A to the shelves: 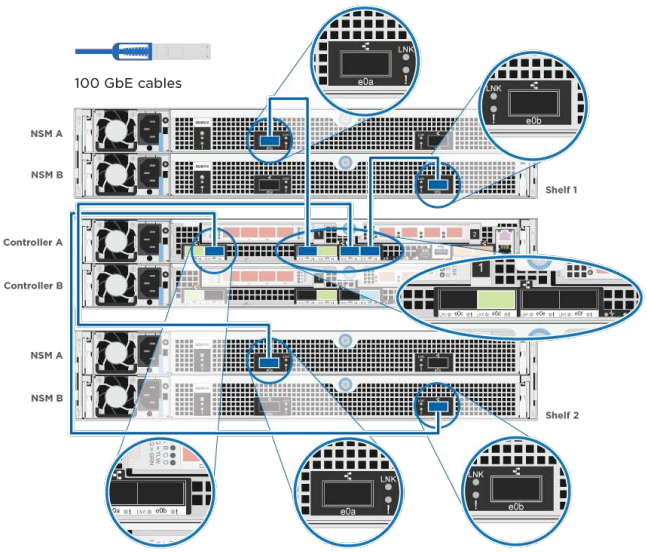

Step	Perform on each controller module
2	Cable controller B to the shelves: 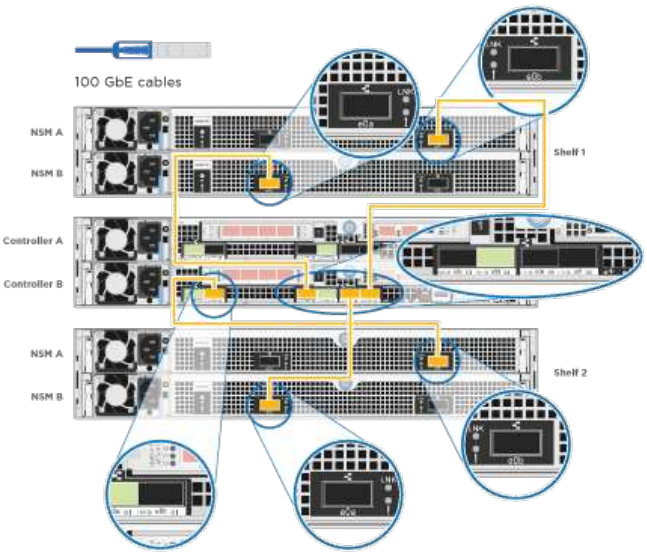

2. To complete setting up your system, see Complete system setup and configuration

Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Completing system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

1. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes

2. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

3. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

4. Select an ONTAP icon listed to discover:

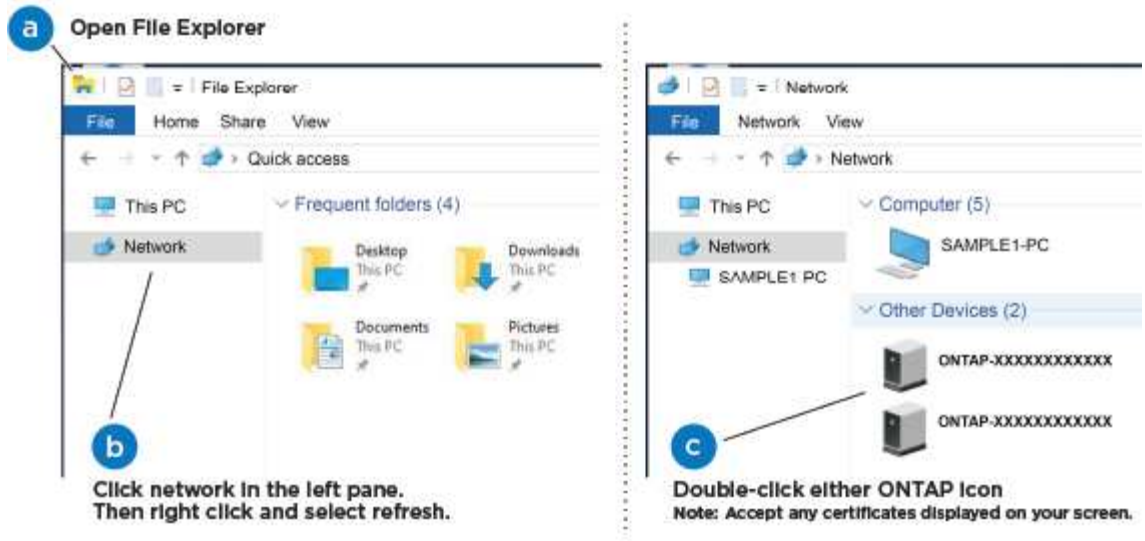

a. Open File Explorer.
b. Click **Network** in the left pane and right-click and select **refresh**.
c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

   XXXXX is the system serial number for the target node.

   System Manager opens.

5. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

   ONTAP Configuration Guide

6. Verify the health of your system by running Config Advisor.
7. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet.

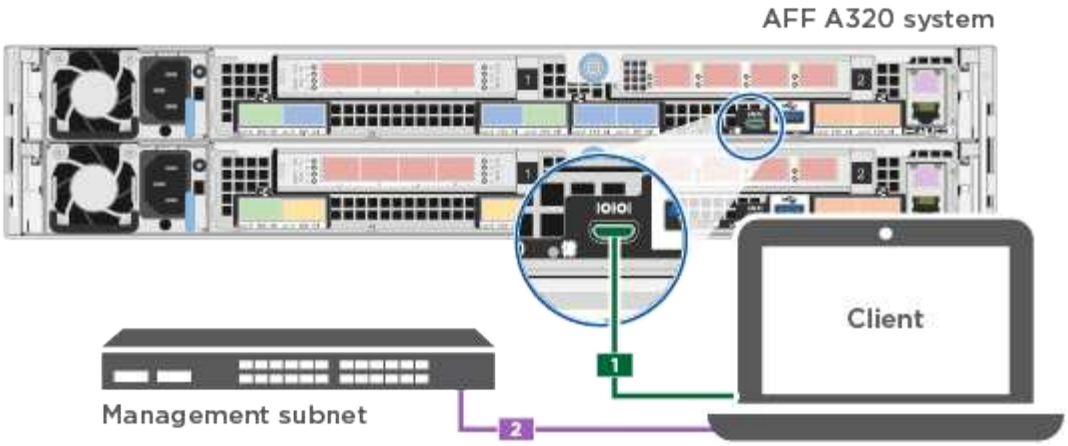

c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Use the following animation to set one or more drive shelf IDs:

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.

   The system begins to boot. Initial booting may take up to eight minutes

4. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:

   a. Point your browser to the node management IP address.

   The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the *NetApp ONTAP Configuration guide*.

   ONTAP Configuration Guide

6. Verify the health of your system by running Config Advisor.

7. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain AFF A320 hardware

For the AFF A320 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Fan

The fan cools the controller.

NVDIMM

```
The NVDIMM (non-volatile dual in-line memory module) manages the data
transfer from the volatile memory to the non-volatile storage, and
maintains data integrity in the event of a power loss or system shutdown.
```

NVDIMM battery

A NVDIMM battery is responsible for maintaining power to the NVDIMM module.

PCIe

A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF A320

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct node:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy node* is the HA partner of the impaired node.

Check encryption key support and status - AFF A320

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:

- If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
- If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the node - AFF A320

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired node. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: System is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - AFF A320

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supply from the power source.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

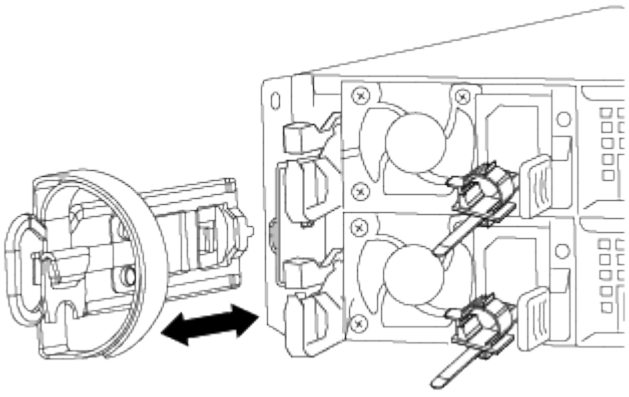

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Remove the controller module from the chassis:

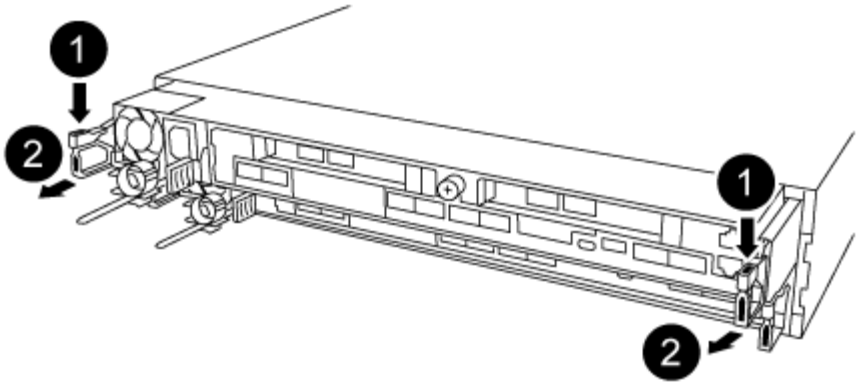

a. Insert your forefinger into the latching mechanism on either side of the controller module.
b. Press down on the orange tab on top of the latching mechanism until it clears the latching pin on the chassis.

   The latching mechanism hook should be nearly vertical and should be clear of the chassis pin.

c. Gently pull the controller module a few inches toward you so that you can grasp the controller module sides.
d. Using both hands, gently pull the controller module out of the chassis and set it on a flat, stable surface.

Step 2: Replace the boot media

You must locate the boot media in the controller module, and then follow the directions to replace it.

1. Open the air duct and locate the boot media using the following illustration or the FRU map on the controller module:
2. Locate and remove the boot media from the controller module:

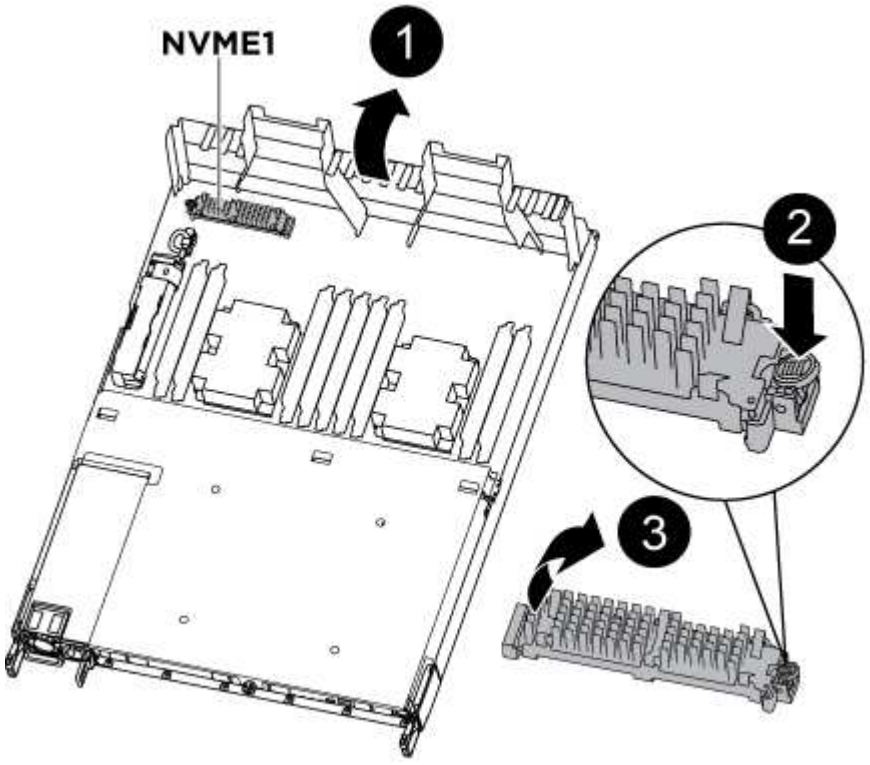

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
b. Rotate the boot media up and gently pull the boot media out of the socket.
    1. Check the boot media to make sure that it is seated squarely and completely in the socket.

       If necessary, remove the boot media and reseat it into the socket.

3. Lock the boot media in place:
    a. Rotate the boot media down toward the motherboard.
    b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
    c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
4. Close the air duct.

Step 3: Transfer the boot image to the boot media using a USB flash drive

The replacement boot media that you installed does not have a boot image, so you need to transfer a boot image using a USB flash drive.

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
    - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
    - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
    a. Download the service image to your work space on your laptop.
    b. Unzip the service image.

       If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

       There are two folders in the unzipped service image file:

        - boot
        - efi

    c. Copy the efi folder to the top directory on the USB flash drive.

       If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

       The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what

the impaired controller is running.

d. Remove the USB flash drive from your laptop.

2. If you have not already done so, close the air duct.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.

5. Plug the power cable into the power supply and reinstall the power cable retainer.
6. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

7. Complete the reinstallation of the controller module:
   a. Make sure the latch arms are locked in the extended position.
   b. Using the latch arms, push the controller module into the chassis bay until it stops.

      

      Do not push down on the latching mechanism at the top of the latch arms. Doing so with raise the locking mechanism and prohibit sliding the controller module into the chassis.

   c. Press down and hold the orange tabs on top of the latching mechanism.
   d. Gently push the controller module into the chassis bay until it is flush with the edges of the chassis.

      

      The latching mechanism arms slide into the chassis.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   e. Release the latches to lock the controller module into place.
   f. If you have not already done so, reinstall the cable management device.
8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the node to boot to LOADER.

9. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

10. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
11. After the image is installed, start the restoration process:
    a. Record the IP address of the impaired node that is displayed on the screen.
    b. Press `y` when prompted to restore the backup configuration.
    c. Press `y` when prompted to overwrite /etc/ssh/ssh_host_dsa_key.

12. From the partner node in advanced privilege level, start the configuration synchronization using the IP address recorded in the previous step: `system node restore-backup -node local -target -address impaired_node_IP_address`
13. If the restore is successful, press `y` on the impaired node when prompted to use the restored copy?.
14. Press `y` when you see confirm backup procedure was successful, and then press `y` when prompted to reboot the node.
15. Verify that the environmental variables are set as expected.
    a. Take the node to the LOADER prompt.

       From the ONTAP prompt, you can issue the command system node halt -skip-lif-migration-before -shutdown true -ignore-quorum-warnings true -inhibit-takeover true.

    b. Check the environment variable settings with the `printenv` command.
    c. If an environment variable is not set as expected, modify it with the `setenv environment-variable-name changed-value` command.
    d. Save your changes using the `savenv` command.
    e. Reboot the node.
16. With the rebooted impaired node displaying the `Waiting for giveback…` message, perform a giveback from the healthy node:

If your system is in…	Then…
An HA pair	After the impaired node is displaying the `Waiting for giveback…` message, perform a giveback from the healthy node: a. From the healthy node: `storage failover giveback -ofnode partner_node_name` The impaired node takes back its storage, finishes booting, and then reboots and is again taken over by the healthy node. ⓘ If the giveback is vetoed, you can consider overriding the vetoes. HA pair management b. Monitor the progress of the giveback operation by using the `storage failover show-giveback` command. c. After the giveback operation is complete, confirm that the HA pair is healthy and that takeover is possible by using the `storage failover show` command. d. Restore automatic giveback if you disabled it using the storage failover modify command.

17. Exit advanced privilege level on the healthy node.

Boot the recovery image - AFF A320

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy node to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address impaired_node_IP_address` d. Return the node to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the node.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press **`y`**.

If your system has…	Then…

No network connection and is in a MetroCluster IP configuration	a. Press `n` when prompted to restore the backup configuration.
	b. Reboot the system when prompted by the system.
	c. Wait for the iSCSI storage connections to connect.
	You can proceed after you see the following messages:
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_auxiliary, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_partner, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_auxiliary, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_partner, address: ip-address).`
	d. Select the **Update flash from backup config** (sync flash) option from the displayed menu.
	If you are prompted to continue with the update, press **`y`**.

4. Ensure that the environmental variables are set as expected:
   a. Take the node to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment_variable_name changed_value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:
   - If your system has onboard keymanager, NSE or NVE configured, go to Post boot media replacement steps for OKM, NSE, and NVE

- If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner node. b. Confirm the target node is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner node.
8. Give back the node using the `storage failover giveback -fromnode local` command
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired node and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore encryption - AFF A320

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF A320

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF A320

To replace the chassis, you must move the fans and controller modules from the impaired chassis to the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the controller modules to the new chassis, and that the chassis is a new component from NetApp.

- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A320

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace
   chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

Replace hardware - AFF A320

Move the fans, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Remove the controller module from the chassis:
   a. Insert your forefinger into the latching mechanism on either side of the controller module.
   b. Press down on the orange tab on top of the latching mechanism until it clears the latching pin on the chassis.

      The latching mechanism hook should be nearly vertical and should be clear of the chassis pin.

   c. Gently pull the controller module a few inches toward you so that you can grasp the controller module sides.
   d. Using both hands, gently pull the controller module out of the chassis and set it on a flat, stable surface.
6. Repeat these steps for the other controller module in the chassis.

Step 2: Move the fans

To move the fan modules to the replacement chassis when replacing the chassis, you must perform a specific sequence of tasks.

1. If you are not already grounded, properly ground yourself.

2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.

4. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

5. Set the fan module aside.
6. Repeat the preceding steps for any remaining fan modules.
7. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The fan LED should be green after the fan is seated and has spun up to operational speed.

10. Repeat these steps for the remaining fan modules.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you must boot your system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Plug the power cables into the power supplies and reinstall the power cable retainers.
4. Complete the reinstallation of the controller module:
   a. Make sure the latch arms are locked in the extended position.
   b. Using the latch arms, push the controller module into the chassis bay until it stops.
   c. Press down and hold the orange tabs on top of the latching mechanism.
   d. Gently push the controller module into the chassis bay until it is flush with the edges of the chassis.

      The latching mechanism arms slide into the chassis.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   e. Release the latches to lock the controller module into place.
   f. Recable the power supply.
   g. If you have not already done so, reinstall the cable management device.
   h. Interrupt the normal boot process by pressing `Ctrl-C`.
5. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - AFF A320

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mccip`

- `non-ha`

b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.
4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller module

Overview of controller module replacement - AFF A320

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - AFF A320

To shut down the impaired controller, you must determine the status of the controller and,

if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - AFF A320

To replace the controller module hardware, you must remove the impaired controller,

move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following images or the written steps to remove the controller module from the chassis.

The following image shows removing the cables and cable management arms from the impaired controller module:

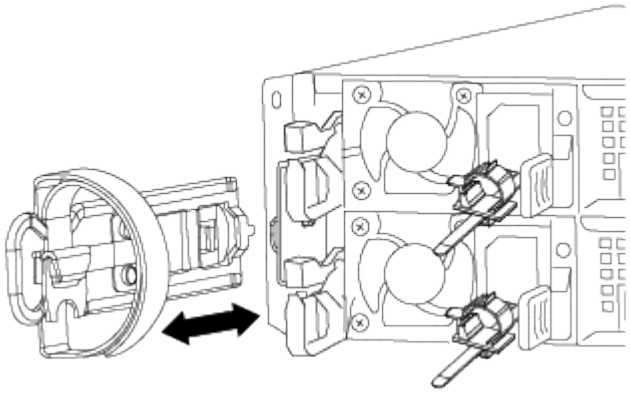

The following image shows removing the impaired controller module from the chassis:

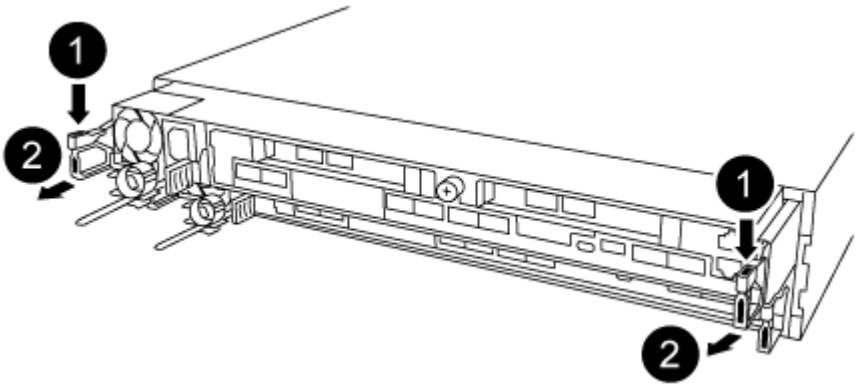

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supply from the power source.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Remove the controller module from the chassis:
   a. Insert your forefinger into the latching mechanism on either side of the controller module.
   b. Press down on the orange tab on top of the latching mechanism until it clears the latching pin on the chassis.

The latching mechanism hook should be nearly vertical and should be clear of the chassis pin.

c. Gently pull the controller module a few inches toward you so that you can grasp the controller module sides.

d. Using both hands, gently pull the controller module out of the chassis and set it on a flat, stable surface.

Step 2: Move the power supplies

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

1. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the blue locking tab.

   The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the NVDIMM battery

To move the NVDIMM battery from the impaired controller module to the replacement controller module, you must perform a specific sequence of steps.

You can use the following illustration or the written steps to move the NVDIMM battery from the impaired controller module to the replacement controller module.

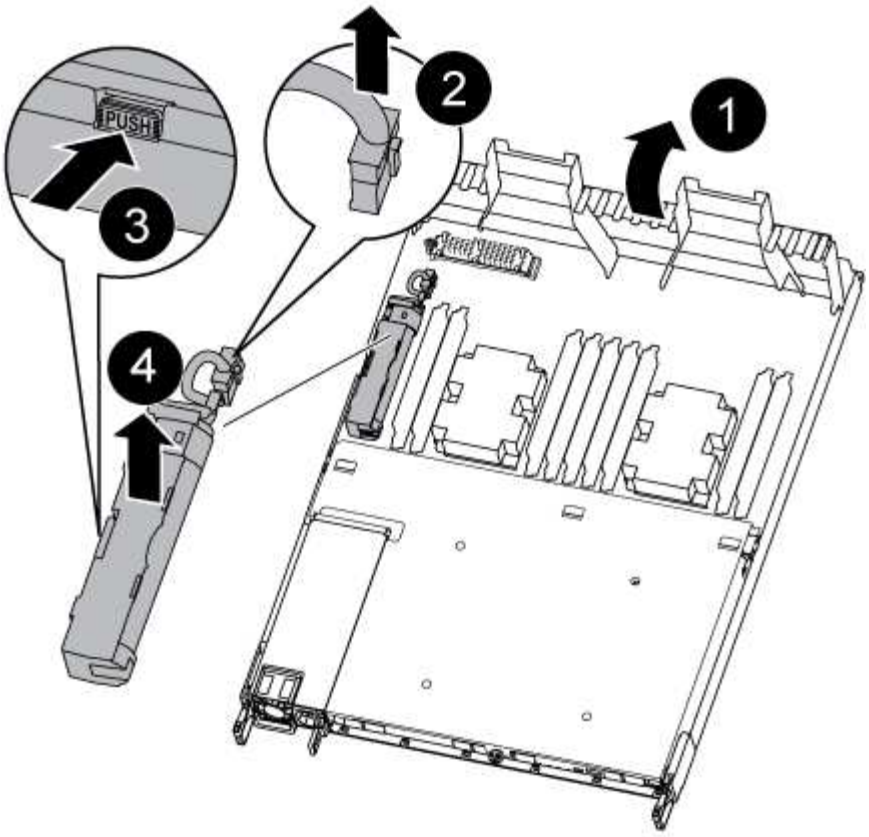

1. Locate the NVDIMM battery in the controller module.
2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
4. Move the battery to the replacement controller module.
5. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.

> Do not plug the battery cable back into the motherboard until instructed to do so.

Step 4: Move the boot media

You must locate the boot media, and then follow the directions to remove it from the impaired controller module and insert it into the replacement controller module.

You can use the following illustration or the written steps to move the boot media from the impaired controller module to the replacement controller module.

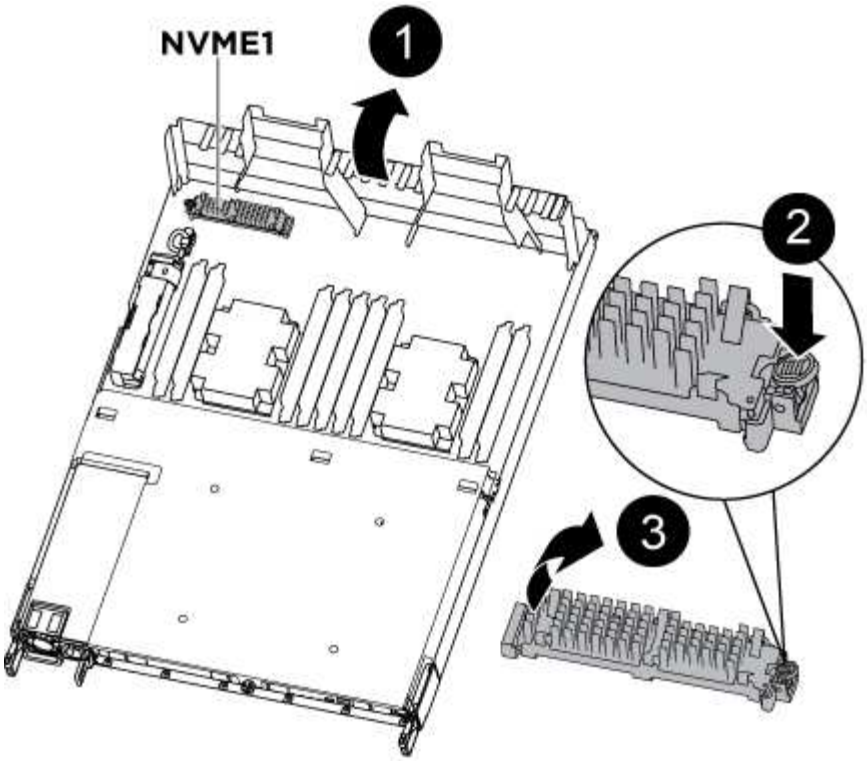

1. Open the air duct and locate the boot media using the following illustration or the FRU map on the controller module:
2. Locate and remove the boot media from the controller module:
   a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
   b. Rotate the boot media up and gently pull the boot media out of the socket.
3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Lock the boot media in place:
   a. Rotate the boot media down toward the motherboard.
   b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
   c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.

Step 5: Move the DIMMs

You need to locate the DIMMs, and then move them from the impaired controller module to the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

You can use the following illustrations or the written steps to move the DIMMs from the impaired controller module to the replacement controller module.

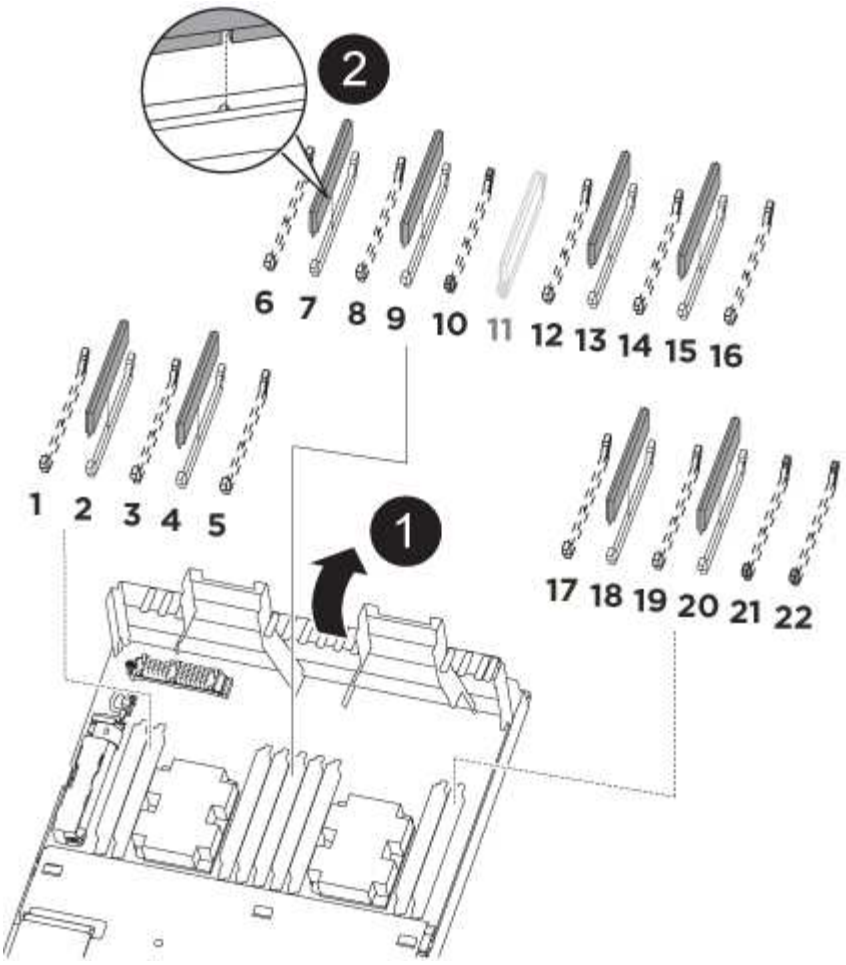

1. Locate the DIMMs on your controller module.

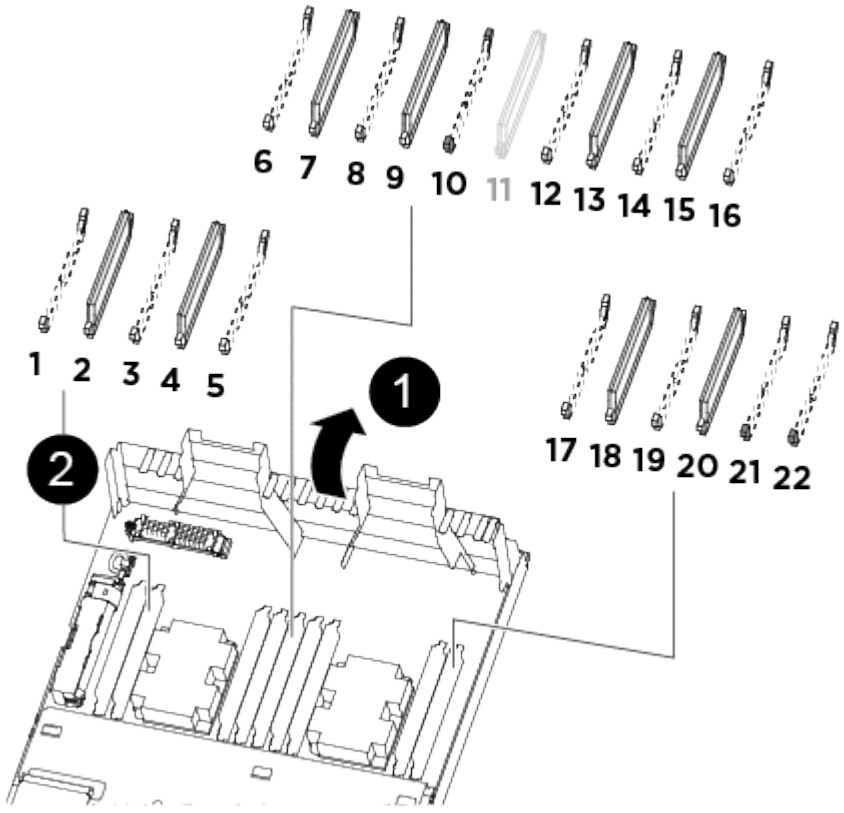

1	Air duct
2	• System DIMMs slots: 2,4, 7, 9, 13, 15, 18, and 20 • NVDIMM slot: 11 The NVDIMM looks significantly different than system DIMMs.

2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Verify that the NVDIMM battery is not plugged into the new controller module.
4. Move the DIMMs from the impaired controller module to the replacement controller module:

   Make sure that you install the each DIMM into the same slot it occupied in the impaired controller module.

   a. Eject the DIMM from its slot by slowly pushing apart the DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

      Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   b. Locate the corresponding DIMM slot on the replacement controller module.
   c. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the DIMM squarely into the socket.

      The DIMMs fit tightly in the socket, but should go in easily. If not, realign the DIMM with the socket and reinsert it.

   d. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.
   e. Repeat these substeps for the remaining DIMMs.
5. Plug the NVDIMM battery into the motherboard.

   Make sure that the plug locks down onto the controller module.

Step 6: Move the PCIe risers

You must move the PCIe risers, with the PCIe cards installed in them, from the impaired controller module to the replacement controller module.

You can use the following illustration or the written steps to move the PCIe risers from the impaired controller module to the replacement controller module.

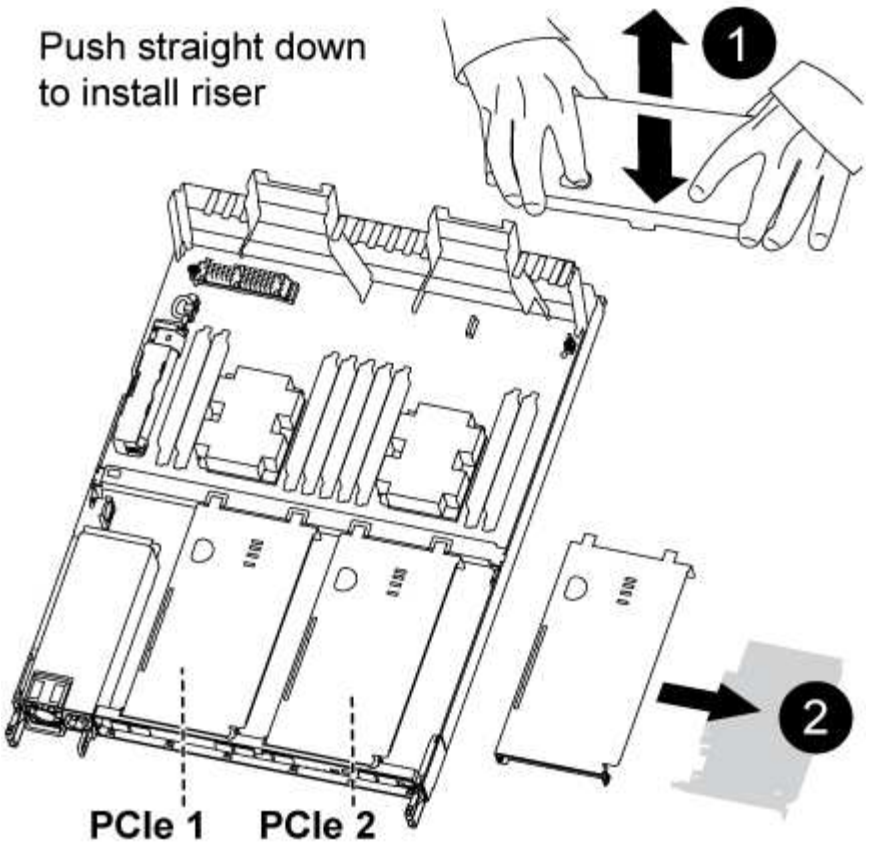

1. Remove the cover over the PCIe risers by unscrewing the blue thumbscrew on the cover, slide the cover toward you, rotate the cover upward, lift it off the controller module, and then set it aside.
2. Remove the empty risers from the replacement controller module.
   a. Place your forefinger into the hole on the left side of the riser module and grasp the riser with your thumb.
   b. Lift the riser straight up and out of the bay, and then set it aside.
   c. Repeat these substeps for the second riser.
3. Move the PCIe risers from the impaired controller module to the same riser bays on the replacement controller module:
   a. Remove a riser from the impaired controller module and move it to the replacement controller module.
   b. Lower the riser straight into the bay, so that it is square with the bay and the pins of the riser slide into the guide holes at the rear of the bay.
   c. Seat the riser into the motherboard socket straight down into the socket by applying even downward pressure along the edges of the riser until it seats.

      The riser should seat smoothly with little resistence. Reseat the riser in the bay if you encounter significant resistance seating the riser into the socket.

   d. Repeat these substeps for the second riser.
   e. Reinstall the cover over the PCIe risers.

Step 7: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller

module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

You can use the following illustration or the written steps to install the replacement controller module in the chassis.

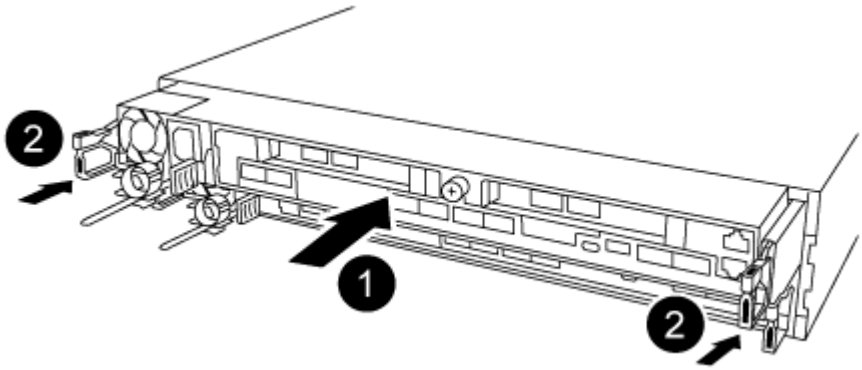

1. If you have not already done so, close the air duct at the rear of the controller module and reinstall the cover over the PCIe cards.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:
   a. Make sure the latch arms are locked in the extended position.
   b. Using the latch arms, push the controller module into the chassis bay until it stops.
   c. Press down and hold the orange tabs on top of the latching mechanism.
   d. Gently push the controller module into the chassis bay until it is flush with the edges of the chassis.

      The latching mechanism arms slide into the chassis.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   e. Release the latches to lock the controller module into place.
   f. Recable the power supply.
   g. If you have not already done so, reinstall the cable management device.
   h. Interrupt the normal boot process by pressing `Ctrl-C`.

Restore and verify the system configuration - AFF A320

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system

settings as necessary.

Step 1: Set and verify the system time after replacing the controller module

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mccip`

- non-ha

3. If the displayed system state of the controller module does not match your system configuration, set the HA state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF A320

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node                Partner         Possible     State Description
------------      ------------      --------
-------------------------------------
node1               node2           false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2               node1           -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed

on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. Verify that the expected volumes are present for each controller: `vol show -node node-name`
9. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF A320

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are

invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A320

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supply from the power source.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

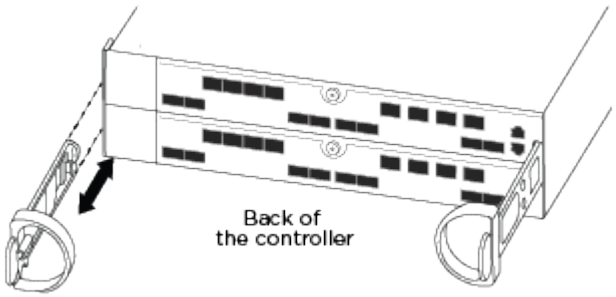

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Remove the controller module from the chassis:

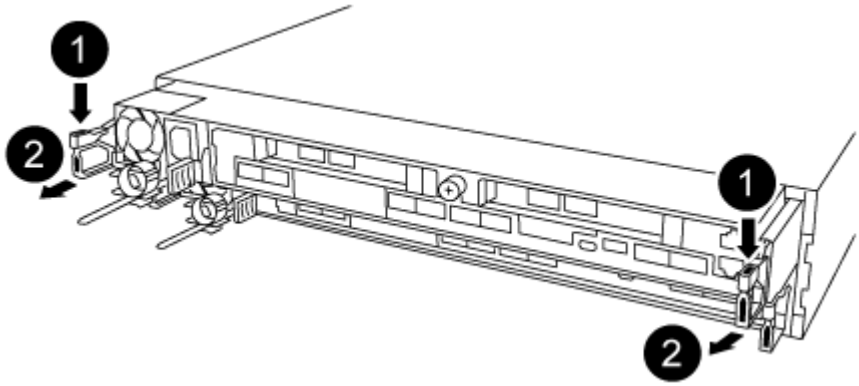

a. Insert your forefinger into the latching mechanism on either side of the controller module.

b. Press down on the orange tab on top of the latching mechanism until it clears the latching pin on the chassis.

   The latching mechanism hook should be nearly vertical and should be clear of the chassis pin.

c. Gently pull the controller module a few inches toward you so that you can grasp the controller module sides.

d. Using both hands, gently pull the controller module out of the chassis and set it on a flat, stable surface.

Step 3: Replace system DIMMs

Replacing a system DIMM involves identifying the target DIMM through the associated error message, locating the target DIMM using the FRU map on the air duct, and then replacing the DIMM.

1. Rotate the air duct to the open position.
2. Locate the DIMMs on your controller module.

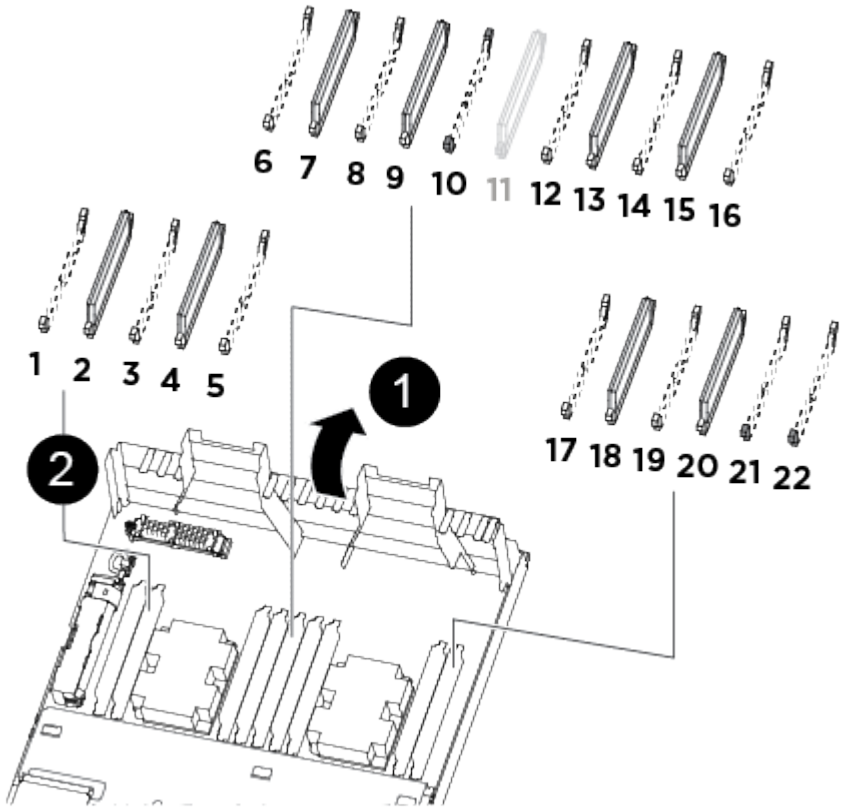

1	Air duct
2	• System DIMMs slots: 2,4, 7, 9, 13, 15, 18, and 20 • NVDIMM slot: 11 The NVDIMM looks significantly different than system DIMMs.

3. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
4. Eject the DIMM from its socket by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the socket.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis.

1. If you have not already done so, close the air duct at the rear of the controller module and reinstall the cover over the PCIe cards.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

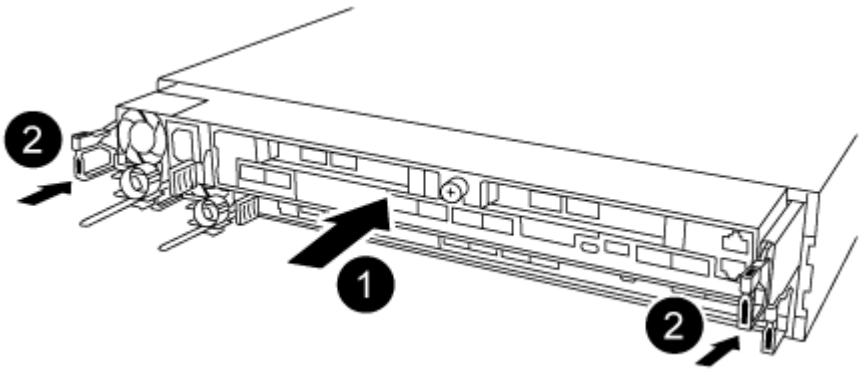

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:
   a. Make sure the latch arms are locked in the extended position.
   b. Using the latch arms, push the controller module into the chassis bay until it stops.
   c. Press down and hold the orange tabs on top of the latching mechanism.
   d. Gently push the controller module into the chassis bay until it is flush with the edges of the chassis.

      The latching mechanism arms slide into the chassis.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   e. Release the latches to lock the controller module into place.
   f. Recable the power supply.
   g. If you have not already done so, reinstall the cable management device.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a fan module - AFF A320

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

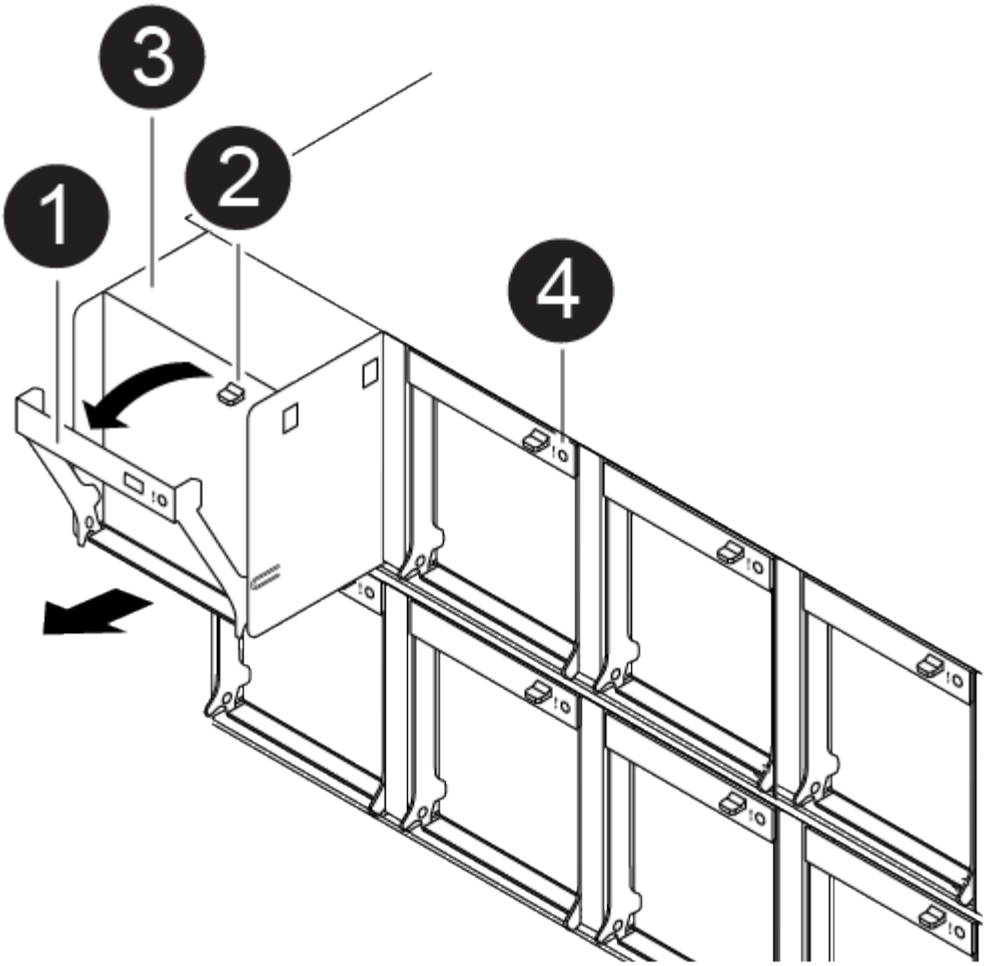

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.

5. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

6. Set the fan module aside.
7. Insert the replacement fan module into the chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The Attention LED should not be lit after the fan is seated and has spun up to operational speed.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Replace an NVDIMM - AFF A320

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supply from the power source.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

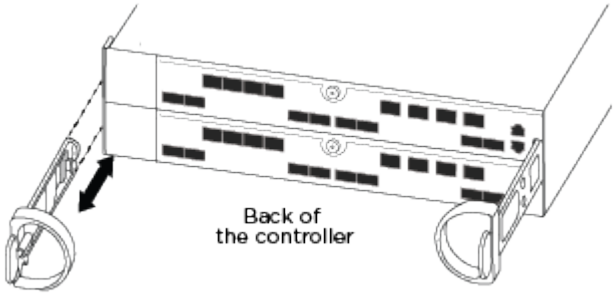

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Remove the controller module from the chassis:

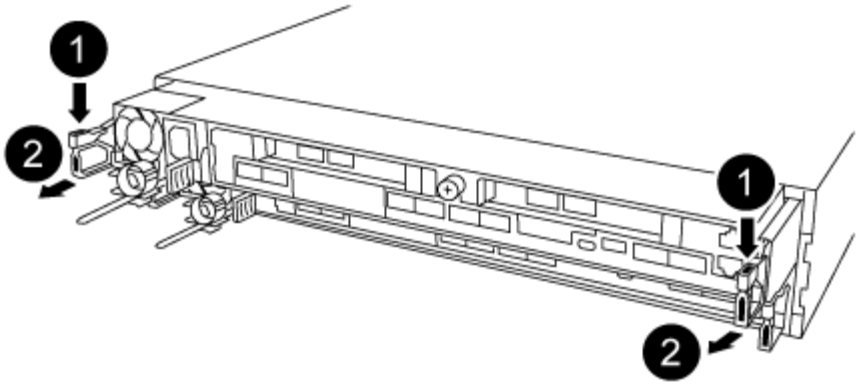

a. Insert your forefinger into the latching mechanism on either side of the controller module.
b. Press down on the orange tab on top of the latching mechanism until it clears the latching pin on the chassis.

   The latching mechanism hook should be nearly vertical and should be clear of the chassis pin.

c. Gently pull the controller module a few inches toward you so that you can grasp the controller module sides.
d. Using both hands, gently pull the controller module out of the chassis and set it on a flat, stable surface.

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the NVDIMM map label on top of the air duct, and then replace it following the specific sequence of steps.

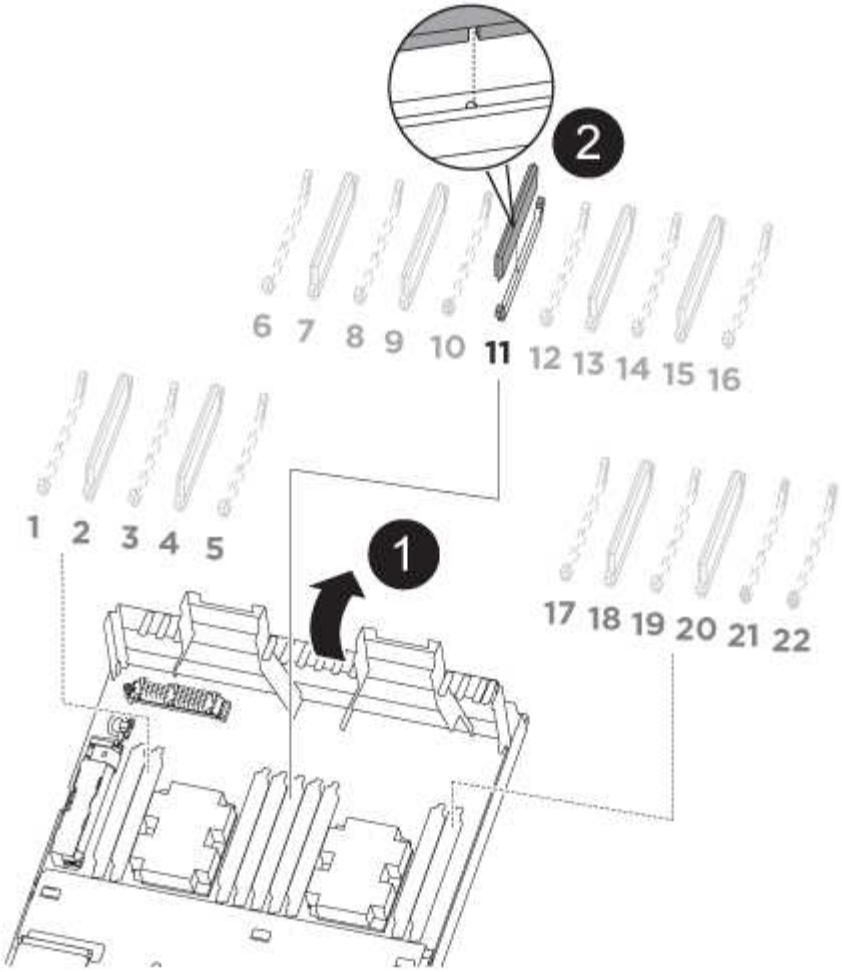

1. Open the air duct and then locate the NVDIMM in slot 11 on your controller module.

   The NVDIMM looks significantly different than system DIMMs.

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

4. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

5. Locate the slot where you are installing the NVDIMM.
6. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
8. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis.

1. If you have not already done so, close the air duct at the rear of the controller module and reinstall the cover over the PCIe cards.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:
   a. Make sure the latch arms are locked in the extended position.
   b. Using the latch arms, push the controller module into the chassis bay until it stops.
   c. Press down and hold the orange tabs on top of the latching mechanism.
   d. Gently push the controller module into the chassis bay until it is flush with the edges of the chassis.

      The latching mechanism arms slide into the chassis.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   e. Release the latches to lock the controller module into place.
   f. Recable the power supply.
   g. If you have not already done so, reinstall the cable management device.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - AFF A320

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supply from the power source.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

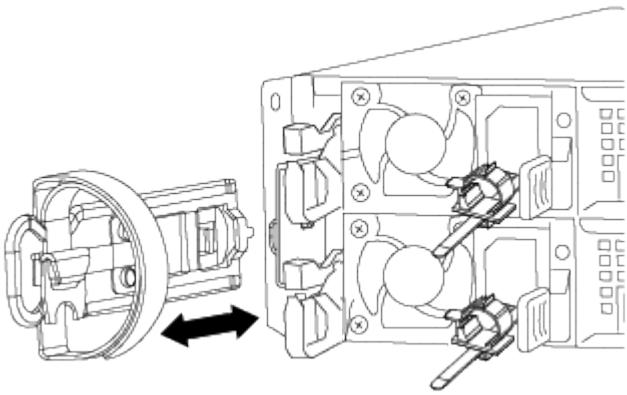

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Remove the controller module from the chassis:

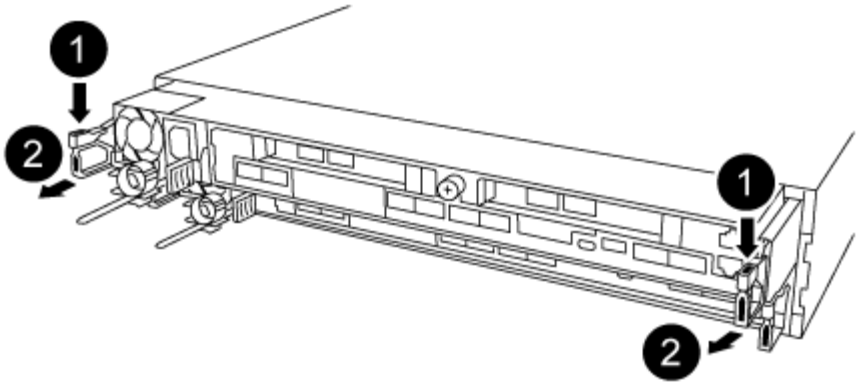

a. Insert your forefinger into the latching mechanism on either side of the controller module.
b. Press down on the orange tab on top of the latching mechanism until it clears the latching pin on the chassis.

   The latching mechanism hook should be nearly vertical and should be clear of the chassis pin.

c. Gently pull the controller module a few inches toward you so that you can grasp the controller module sides.
d. Using both hands, gently pull the controller module out of the chassis and set it on a flat, stable surface.

Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module.

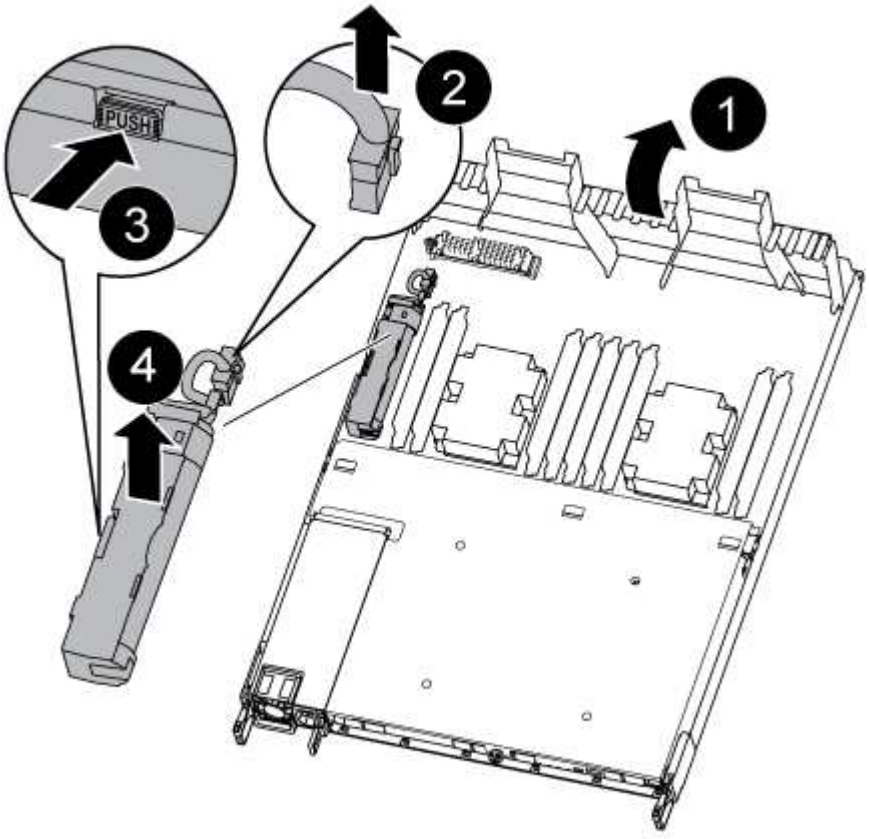

1. Open the air duct and locate the NVDIMM battery.
2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
4. Remove the replacement battery from its package.
5. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.
6. Plug the battery plug back into the controller module, and then close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it..

1. If you have not already done so, close the air duct at the rear of the controller module and reinstall the cover over the PCIe cards.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

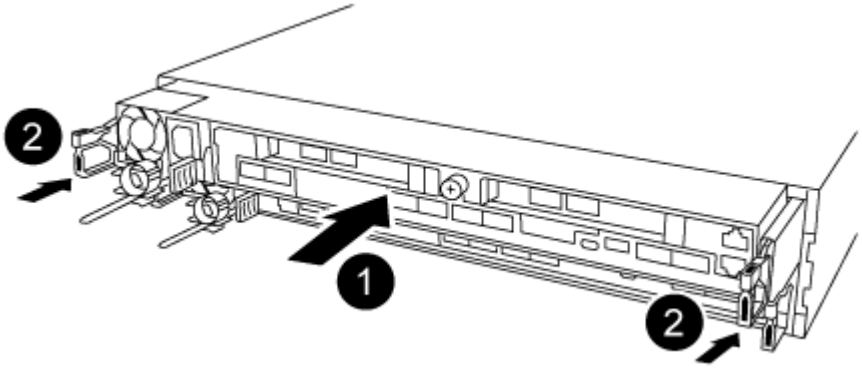

Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:
   a. Make sure the latch arms are locked in the extended position.
   b. Using the latch arms, push the controller module into the chassis bay until it stops.
   c. Press down and hold the orange tabs on top of the latching mechanism.
   d. Gently push the controller module into the chassis bay until it is flush with the edges of the chassis.

The latching mechanism arms slide into the chassis.

The controller module begins to boot as soon as it is fully seated in the chassis.

e. Release the latches to lock the controller module into place.

f. Recable the power supply.

g. If you have not already done so, reinstall the cable management device.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe card - AFF A320

To replace a PCIe card, you must disconnect the cables from the cards, remove the SFP and QSFP modules from the cards before removing the riser, reinstall the riser, and then reinstall the SFP and QSFP modules before cabling the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supply from the power source.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

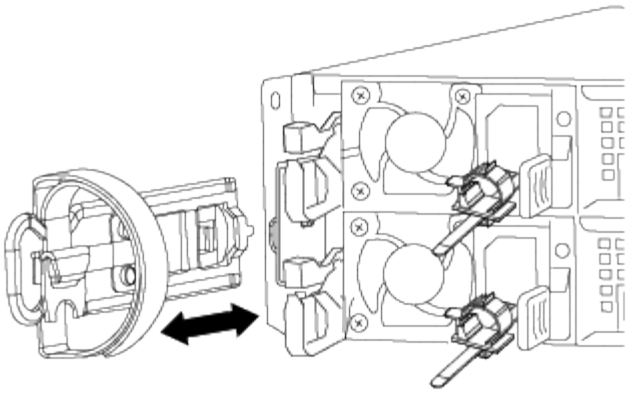

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Remove the controller module from the chassis:

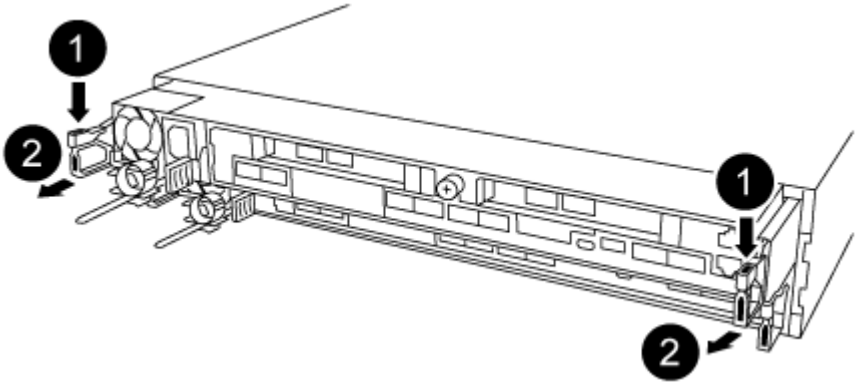

   a. Insert your forefinger into the latching mechanism on either side of the controller module.
   b. Press down on the orange tab on top of the latching mechanism until it clears the latching pin on the chassis.

      The latching mechanism hook should be nearly vertical and should be clear of the chassis pin.

   c. Gently pull the controller module a few inches toward you so that you can grasp the controller module sides.
   d. Using both hands, gently pull the controller module out of the chassis and set it on a flat, stable surface.

Step 3: Replace a PCIe card

You must remove the PCIe riser containing the failed PCIe card from the controller module, remove the failed PCIe card from the riser, install the replacement PCIe card in the riser, and then reinstall the riser into the controller module.

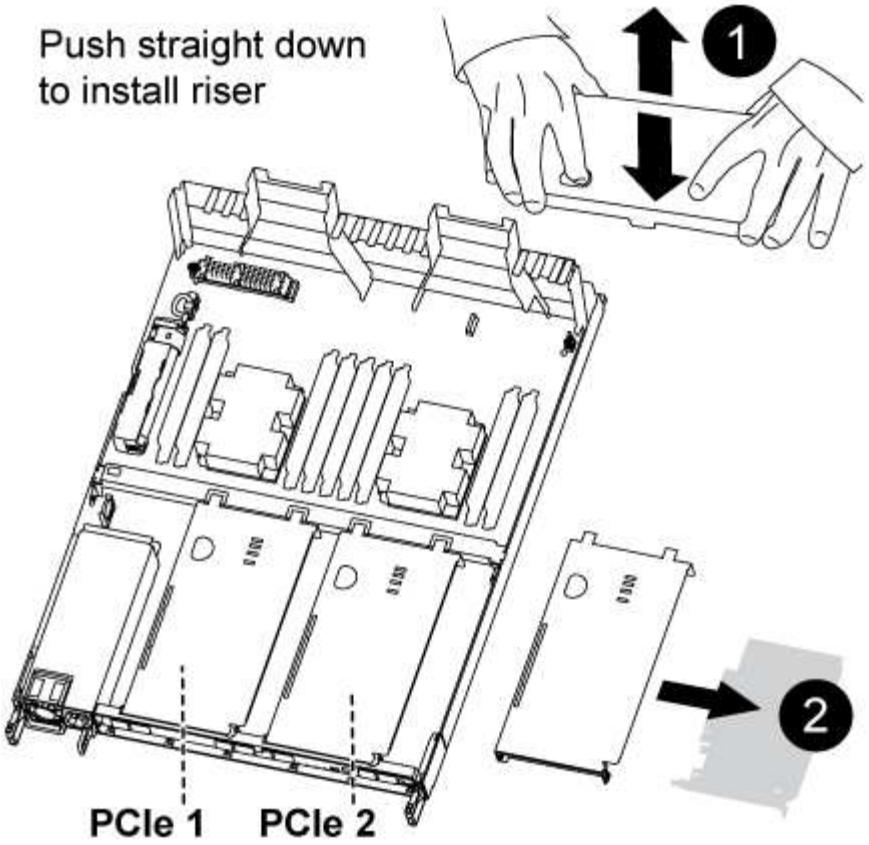

1. Remove the cover over the PCIe risers by unscrewing the blue thumbscrew on the cover, slide the cover toward you, rotate the cover upward, lift it off the controller module, and then set it aside.
2. Remove the riser with the failed PCIe card:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Place your forefinger into the hole on the left side of the riser module and grasp the riser with your thumb.
   c. Lift the riser straight out of the socket and set it aside.
3. Replace the card in the riser:
   a. Place the riser on a stable surface, and then turn the riser so that you can access the PCIe card.
   b. Place your thumbs just below the bottom edge of the PCIe card on either side of the socket, and then gently push up to release the card from the socket.
   c. Slide the card out of the riser and set it aside.
   d. Align the replacement card bezel with the edge of the riser and the outside edge of the card with the alignment guide on the left side of the riser.
   e. Gently slide the card until the card connector aligns with the riser socket, and then gently push the card down into the socket.
4. Reinstall the riser in the controller module:
   a. Align the riser over the opening so that the front edges of the riser are directly over the openings on the riser bay.
   b. Aligning the back edge of the riser so that the pins on the underside of the riser are over the holes in the sheet metal at the back riser bay.
   c. Apply even downward pressure to seat the riser straight down into the socket on the controller module.

d. Reinstall the PCIe riser cover on the controller module.

Sep 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

1. If you have not already done so, close the air duct at the rear of the controller module and reinstall the cover over the PCIe cards.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

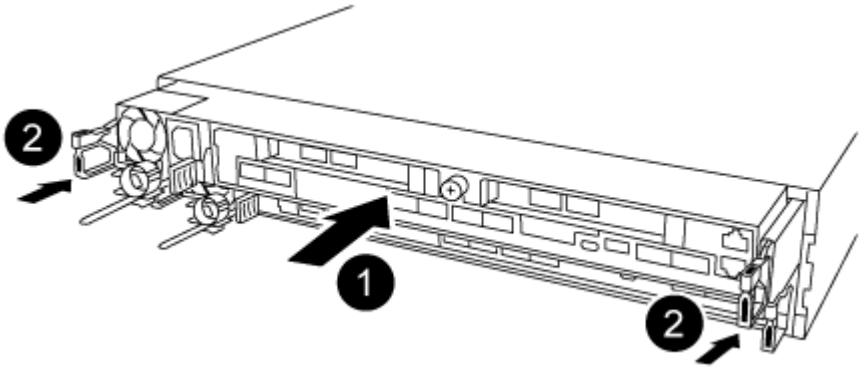

Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module:
   a. Make sure the latch arms are locked in the extended position.
   b. Using the latch arms, push the controller module into the chassis bay until it stops.
   c. Press down and hold the orange tabs on top of the latching mechanism.
   d. Gently push the controller module into the chassis bay until it is flush with the edges of the chassis.

   The latching mechanism arms slide into the chassis.

   The controller module begins to boot as soon as it is fully seated in the chassis.

   e. Release the latches to lock the controller module into place.
   f. Recable the power supply.
   g. If you have not already done so, reinstall the cable management device.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

1. Recable the system, as needed.

If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - AFF A320

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting the replacement PSU to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- Power supplies are auto-ranging.

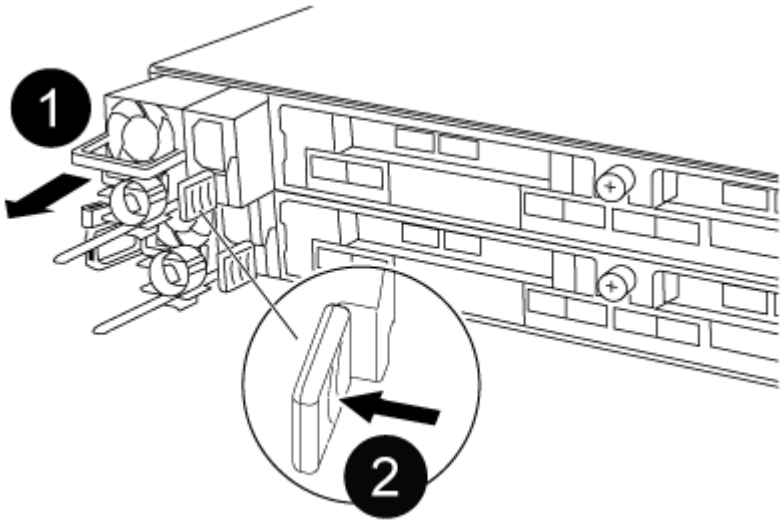

Figure 1. Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
3. Disconnect the power supply:
   a. Open the power cable retainer, and then unplug the power cable from the power supply.

b. Unplug the power cable from the power source.

4. Remove the power supply:
   a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
   b. Press the blue locking tab to release the power supply from the chassis.
   c. Using both hands, pull the power supply out of the chassis, and then set it aside.
5. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

6. Rotate the cam handle so that it is flush against the power supply.
7. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A320

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the

impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the RTC battery

You need to locate the RTC battery inside the controller module, and then follow the specific sequence of steps.

Step 3: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supply from the power source.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

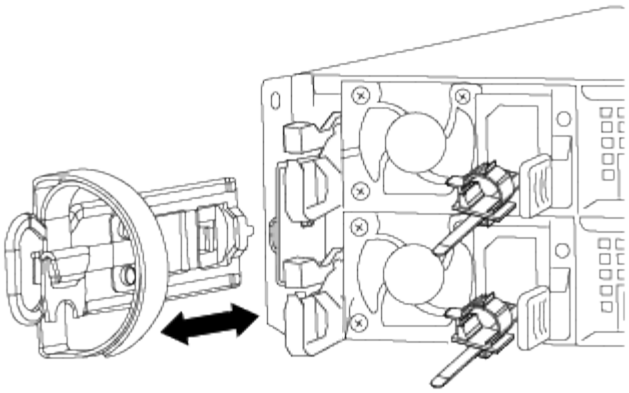

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Remove the controller module from the chassis:

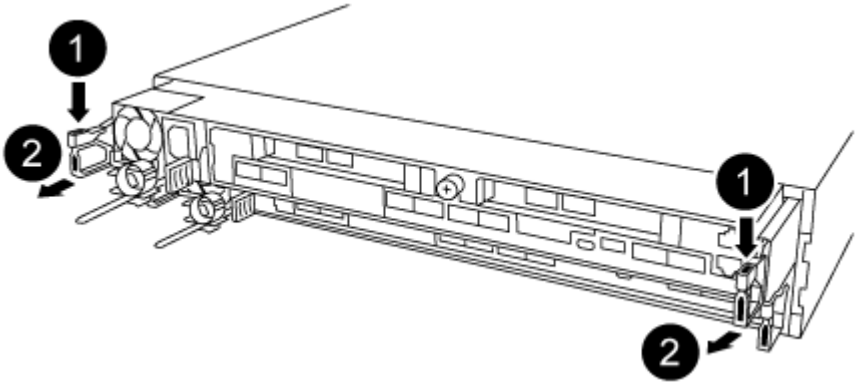

a. Insert your forefinger into the latching mechanism on either side of the controller module.
b. Press down on the orange tab on top of the latching mechanism until it clears the latching pin on the chassis.

   The latching mechanism hook should be nearly vertical and should be clear of the chassis pin.

c. Gently pull the controller module a few inches toward you so that you can grasp the controller module sides.
d. Using both hands, gently pull the controller module out of the chassis and set it on a flat, stable surface.

Step 4: Replace the RTC battery

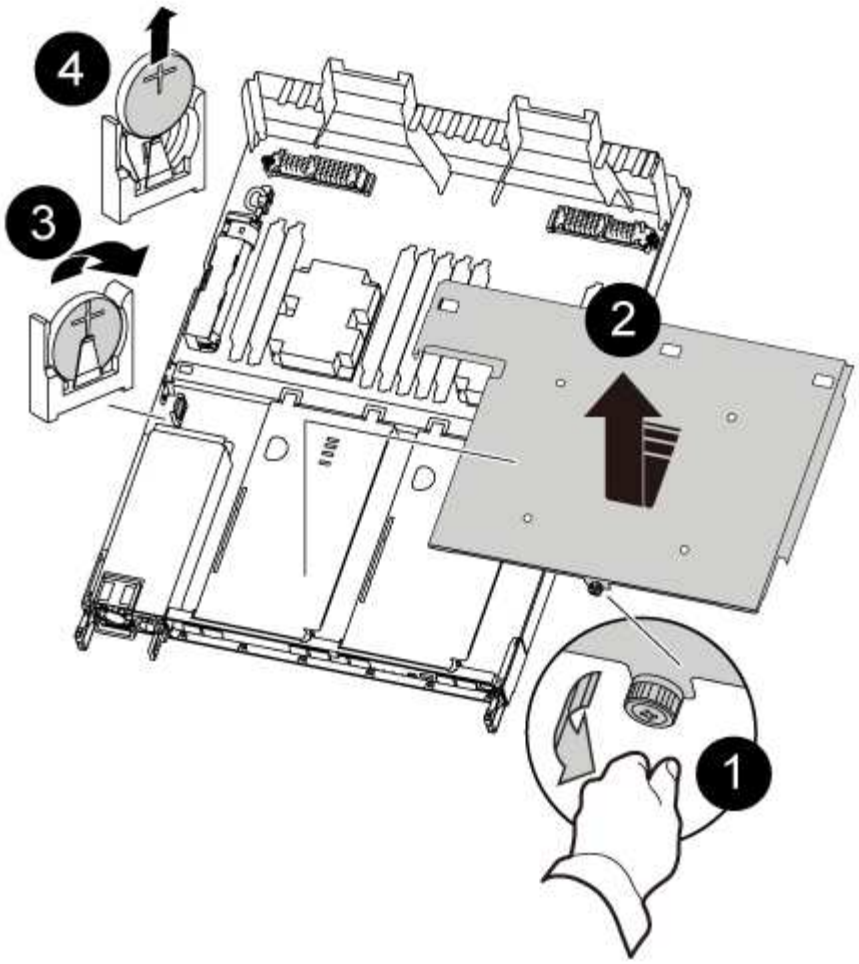

1. Remove the PCIe cover.
   a. Unscrew the blue thumbscrew located above the onboard ports at the back of the controller module.
   b. Slide the cover toward you and rotate the cover upward.
   c. Remove the cover and set it aside.
2. Locate, remove, and then replace the RTC battery:
   a. Using the FRU map, locate the RTC battery on the controller module.
   b. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

   c. Remove the replacement battery from the antistatic shipping bag.
   d. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
3. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
4. Reinstall the PCIe cover on the controller module.

Step 5: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. Make sure the latch arms are locked in the extended position.
   b. Using the latch arms, push the controller module into the chassis bay until it stops.

      Do not push down on the latching mechanism at the top of the latch arms. Doing so with raise the locking mechanism and prohibit sliding the controller module into the chassis.

   c. Press down and hold the orange tabs on top of the latching mechanism.
   d. Gently push the controller module into the chassis bay until it is flush with the edges of the chassis.

      The latching mechanism arms slide into the chassis.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   e. Release the latches to lock the controller module into place.
   f. If you have not already done so, reinstall the cable management device.
   g. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF A700 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

For MetroCluster configurations, see either:

- Install MetroCluster IP configuration
- Install MetroCluster Fabric-Attached configuration

Quick steps - AFF A700

This guide gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Access the *Installation and Setup Instructions* PDF poster:

AFF A700 Installation and Setup Instructions

FAS9000 Installation and Setup Instructions

Video steps - AFF A700

The following video shows how to install and cable your new system.

Animation - Install and setup of an AFF A700 or FAS9000

Detailed guide - AFF A700

This guide gives detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your system, you need to create an account on the NetApp Support Site, register your system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
10 GbE network cable	X6566B-2-R6, (112-00299), 2m X6566B-3-R6, 112-00300, 3m X6566B-5-R6 , 112-00301, 5m		Network cable
40 GbE network cable 40 GbE cluster interconnect	X66100-1,112-00542, 1m X66100-3,112-00543, 3m		40 GbE network Cluster interconnect

Type of cable…	Part number and length	Connector type	For…
100 GbE network cable 100 GbE storage cable	X66211A-05 (112-00595), 0.5m X66211A-1 (112-00573), 1m X66211A-2 (112-00574), 2m X66211A-5 (112-00574), 5m		Network cable Storage cable ⓘ This cable applies to AFF A700 only.
Optical network cables (order dependent)	X6553-R6 (112-00188), 2m X6536-R6 (112-00090), 5m		FC host network
Cat 6, RJ-45 (order dependent)	Part numbers X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network and Ethernet data
Storage	X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		Storage
Micro-USB console cable	Not applicable		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	Not applicable		Powering up the system

4. Review the *NetApp ONTAP Confiuration Guide* and collect the required information listed in that guide.

   ONTAP Configuration Guide

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   ⓘ You need to be aware of the safety concerns associated with the weight of the system.

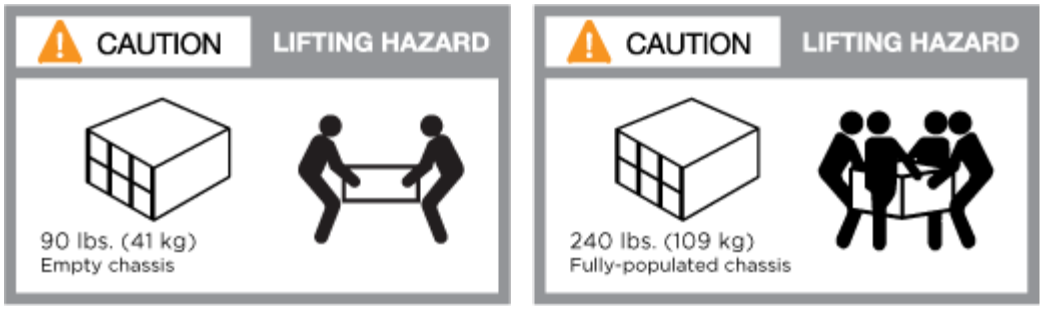

The label on the left indicates an empty chassis, while the label on the right indicates a fully-populated system.

3. Attach cable management devices (as shown).

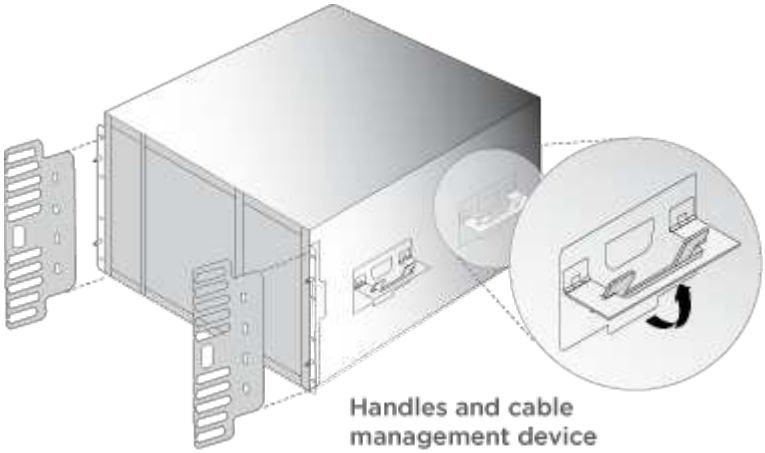

4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Two-node switchless cluster

Management network, data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all networking module ports.

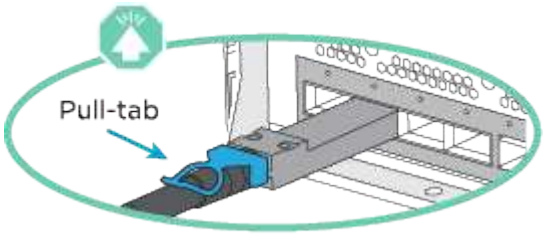

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Cable a two-node switchless cluster

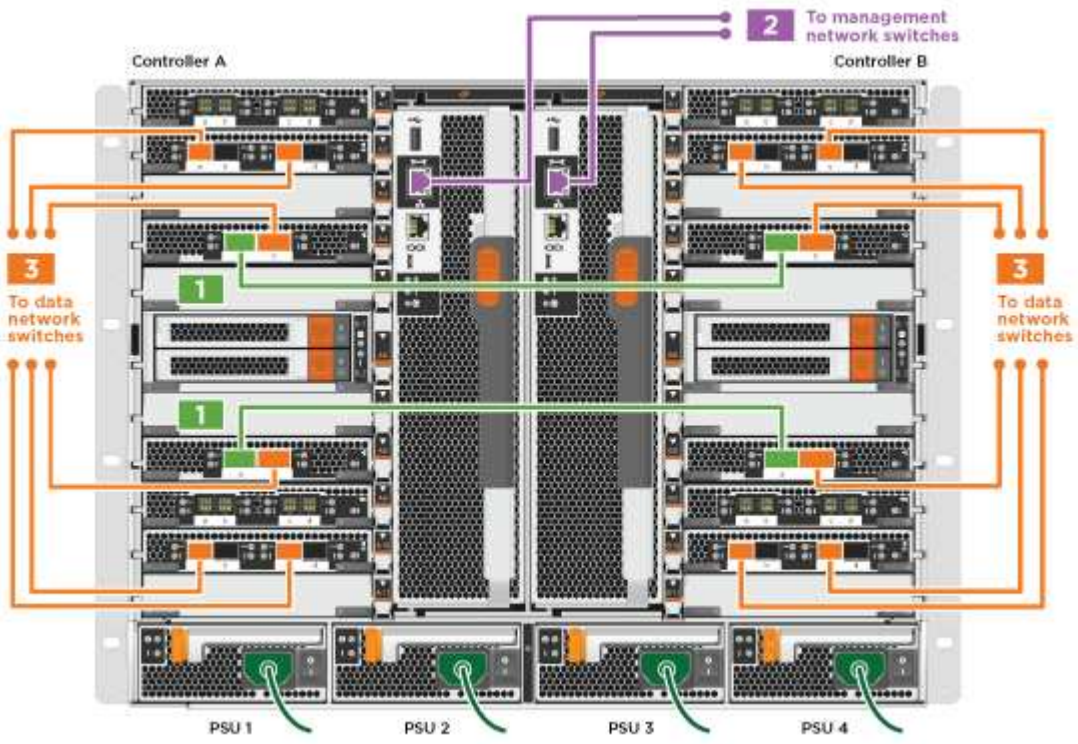

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Option 2: Switched cluster

Management network, data network, and management ports on the controllers are connected to switches. The cluster interconnect and HA ports are cabled on to the cluster/HA switch.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all networking module ports.

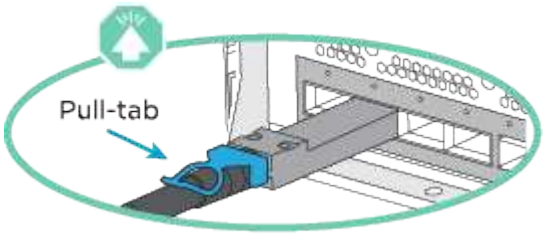

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Switched cluster cabling

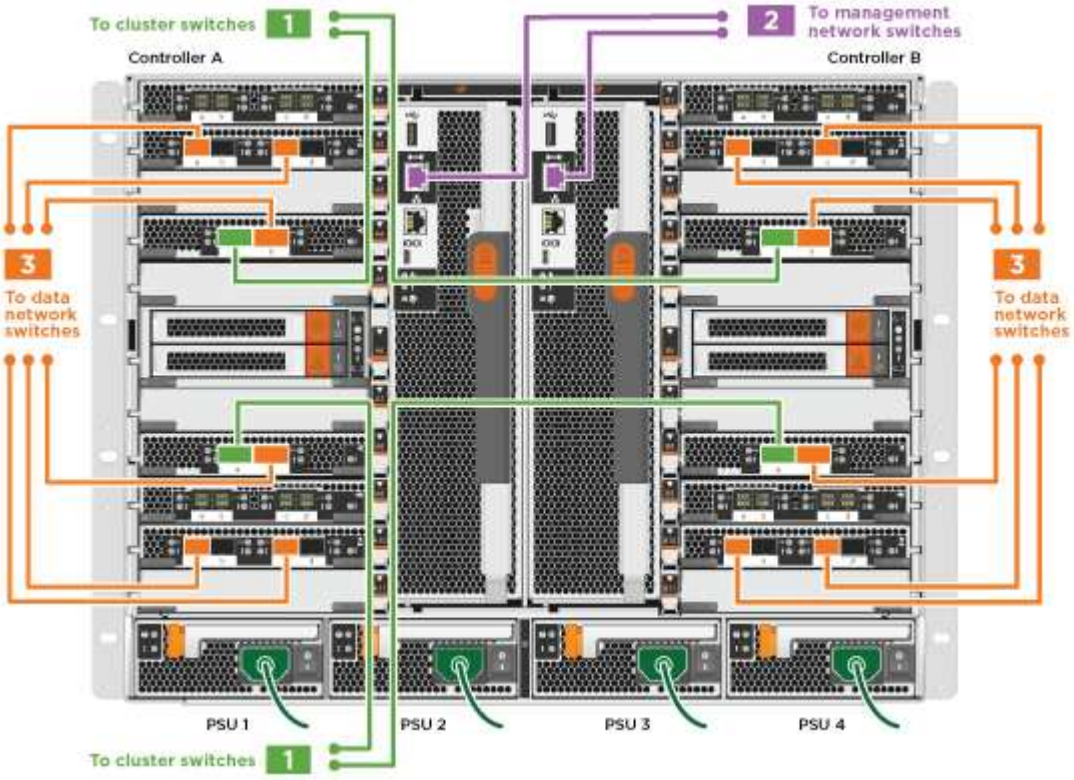

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Step 4: Cable controllers to drive shelves

You can cable your new system to DS212C, DS224C, or NS224 shelves, depending on if it is an AFF or FAS system.

Option 1: Cable the controllers to DS212C or DS224C drive shelves

You must cable the shelf-to-shelf connections, and then cable both controllers to the DS212C or DS224C drive shelves.

The cables are inserted into the drive shelf with the pull-tabs facing down, while the other end of the cable is inserted into the controller storage modules with the pull-tabs up.

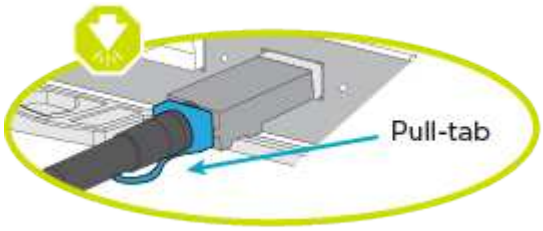

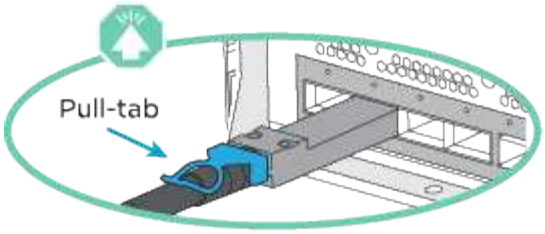

Steps

1. Use the following animations or illustrations to cable your drive shelves to your controllers.

   The examples use DS224C shelves. Cabling is similar with other supported SAS drive shelves.

   - Cabling SAS shelves in FAS9000, AFF A700, and ASA AFF A700, ONTAP 9.7 and earlier:

     Animation - Cable SAS storage - ONTAP 9.7 and earlier

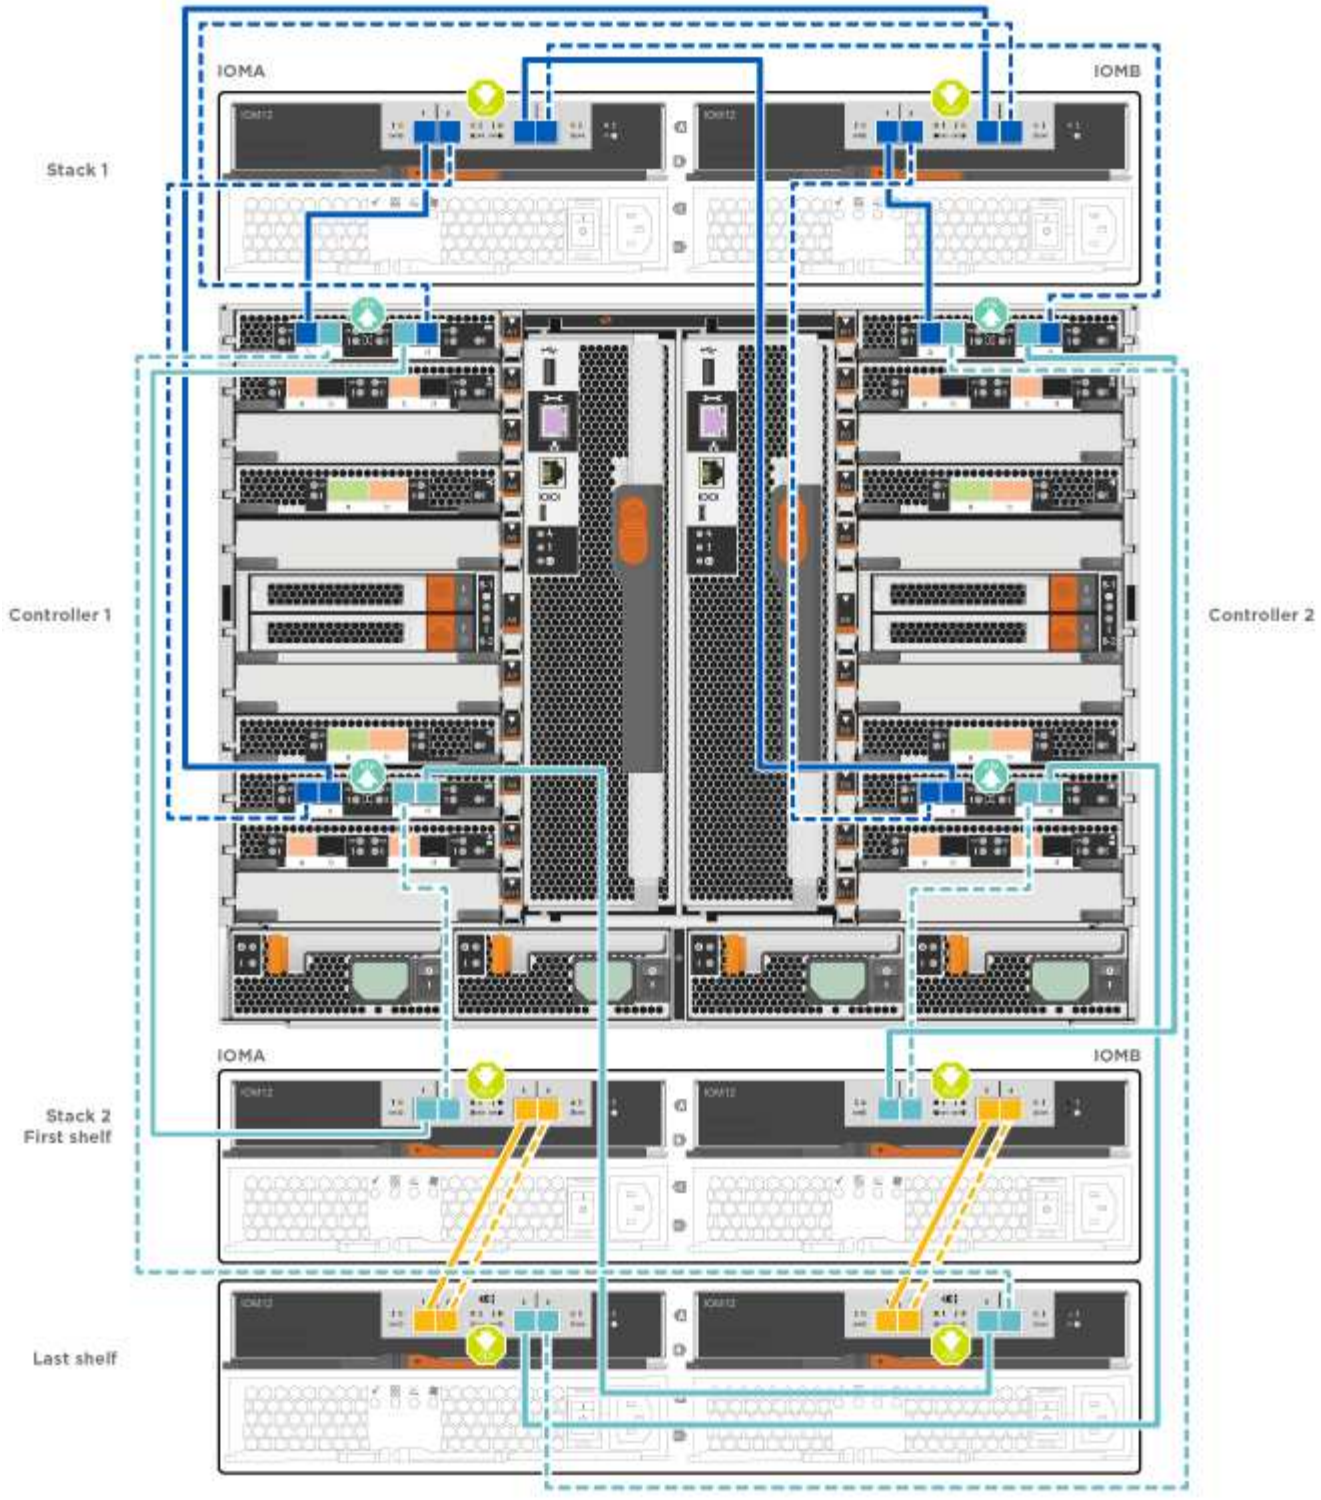

◦ Cabling SAS shelves in FAS9000, AFF A700, and ASA AFF A700, ONTAP 9.8 and later:

Animation - Cable SAS storage - ONTAP 9.8 and later

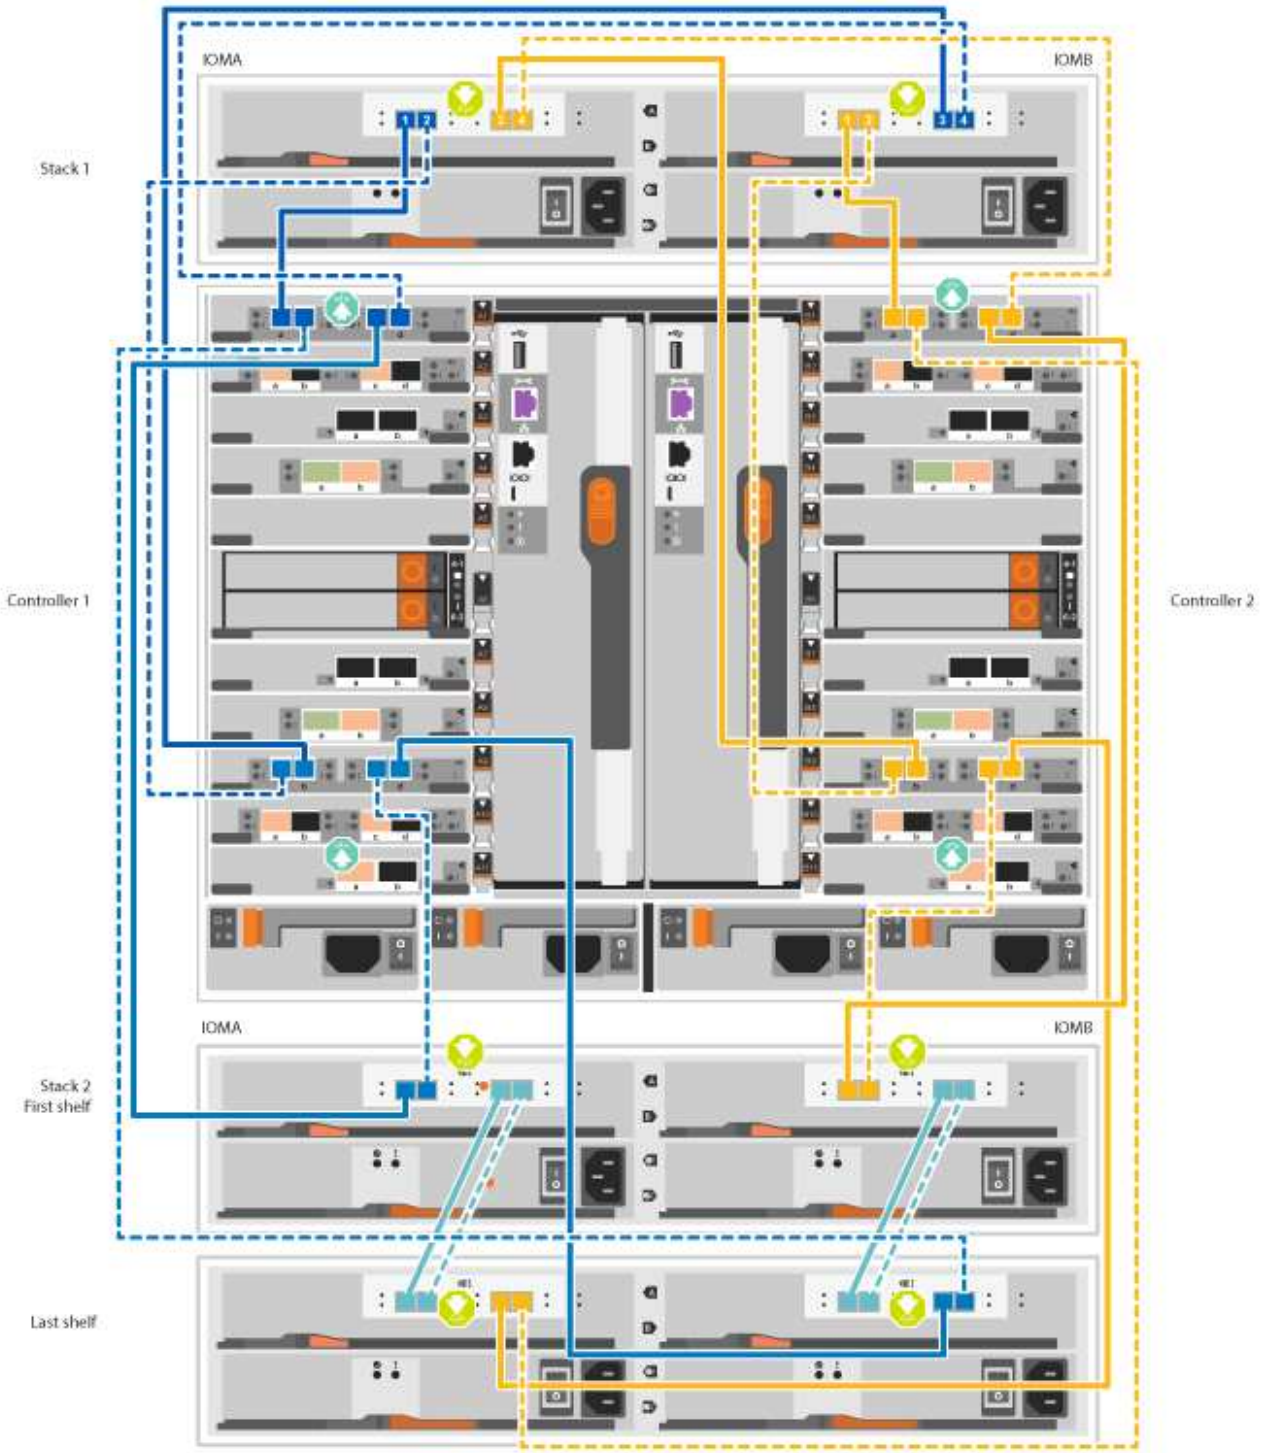

If you have more than one drive shelf stack, see the *Installation and Cabling Guide* for your drive shelf type.

Install and cable shelves for a new system installation - shelves with IOM12 modules

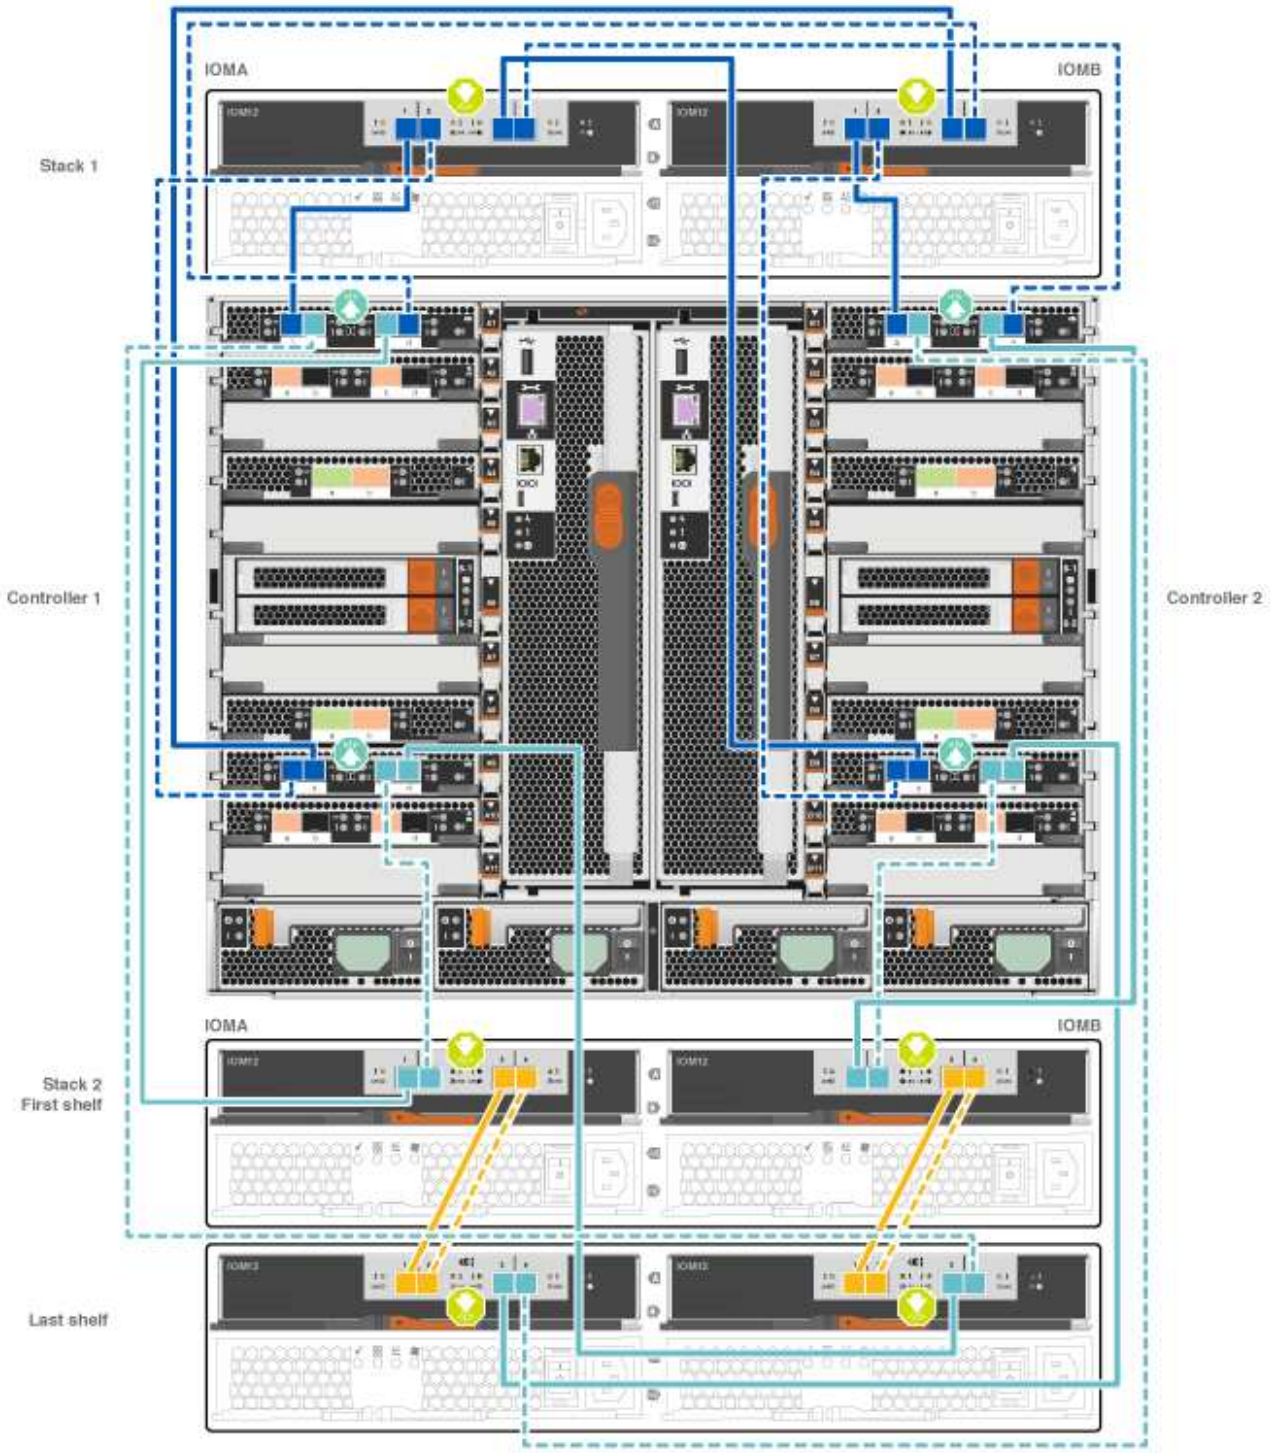

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Option 2: Cable the controllers to a single NS224 drive shelf in AFF A700 and ASA AFF A700 systems running ONTAP 9.8 and later only

You must cable each controller to the NSM modules on the NS224 drive shelf on an AFF A700 or ASA AFF A700 running system ONTAP 9.8 or later.

- This task applies to AFF A700 and ASA AFF A700 running ONTAP 9.8 or later only.
- The systems must have at least one X91148A module installed in slots 3 and/or 7 for each controller. The

animation or illustrations show this module installed in both slots 3 and 7.

- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the storage modules are up, while the pull tabs on the shelves are down.

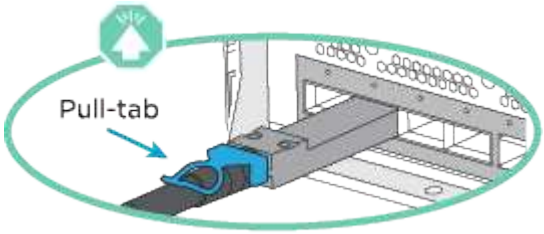

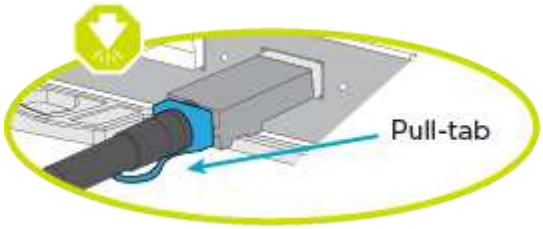

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following animation or illustrations to cable your controllers with two X91148A storage modules to a single NS224 drive shelf, or use the diagram to cable your controllers with one X91148A storage module to a single NS224 drive shelf.

   Animation - Cable a single NS224 shelf - ONTAP 9.8 and later

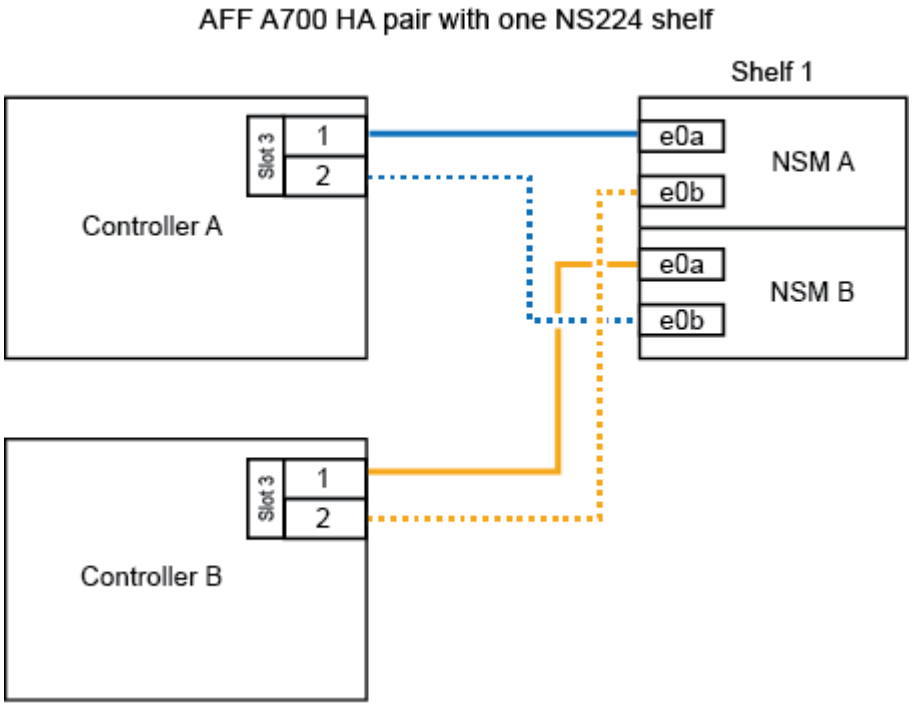

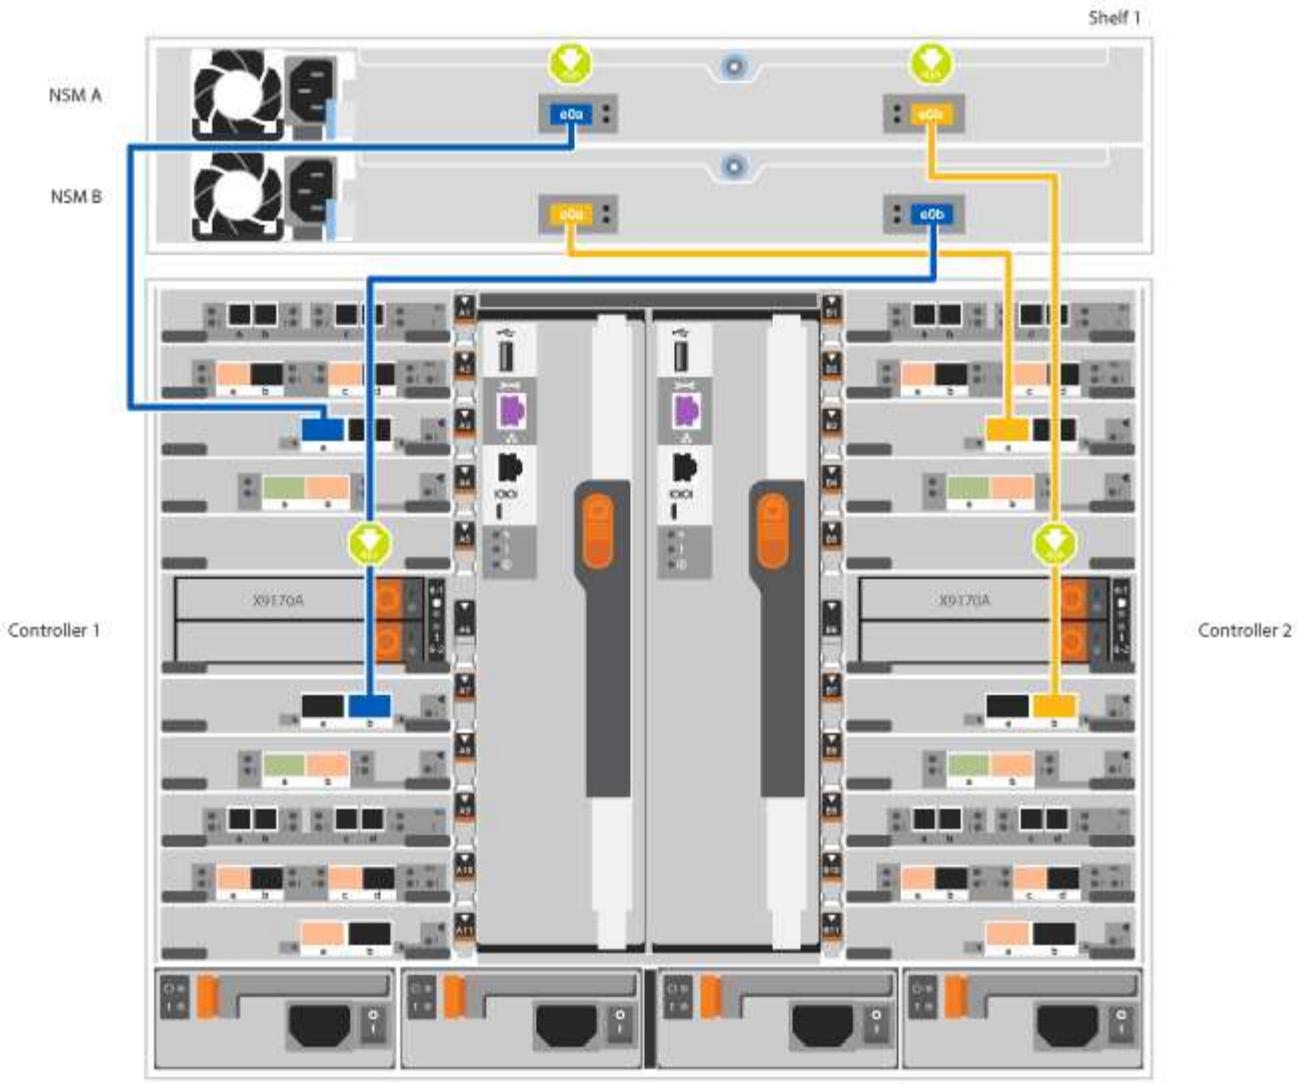

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Option 3: Cable the controllers to two NS224 drive shelves in AFF A700 and ASA AFF A700 systems running ONTAP 9.8 and later only

You must cable each controller to the NSM modules on the NS224 drive shelves on an AFF A700 or ASA AFF A700 running system ONTAP 9.8 or later.

- This task applies to AFF A700 and ASA AFF A700 running ONTAP 9.8 or later only.
- The systems must have two X91148A modules, per controller, installed in slots 3 and 7.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the storage modules are up, while the pull tabs on the shelves are down.

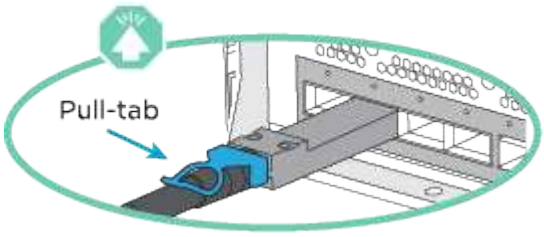

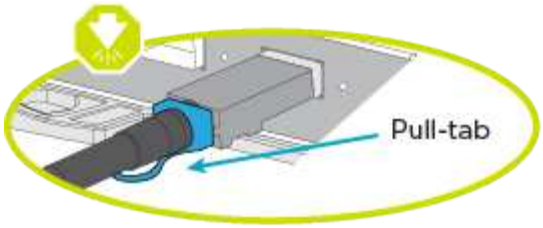

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following animation or illustrations to cable your controllers to two NS224 drive shelves.

Animation - Cable two NS224 shelves - ONTAP 9.8 and later

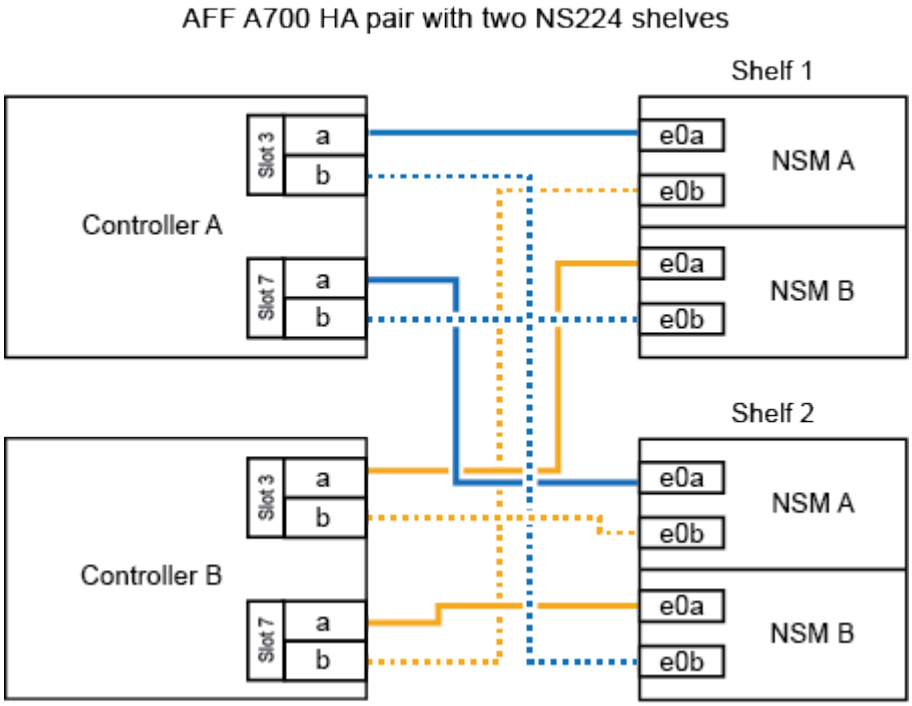

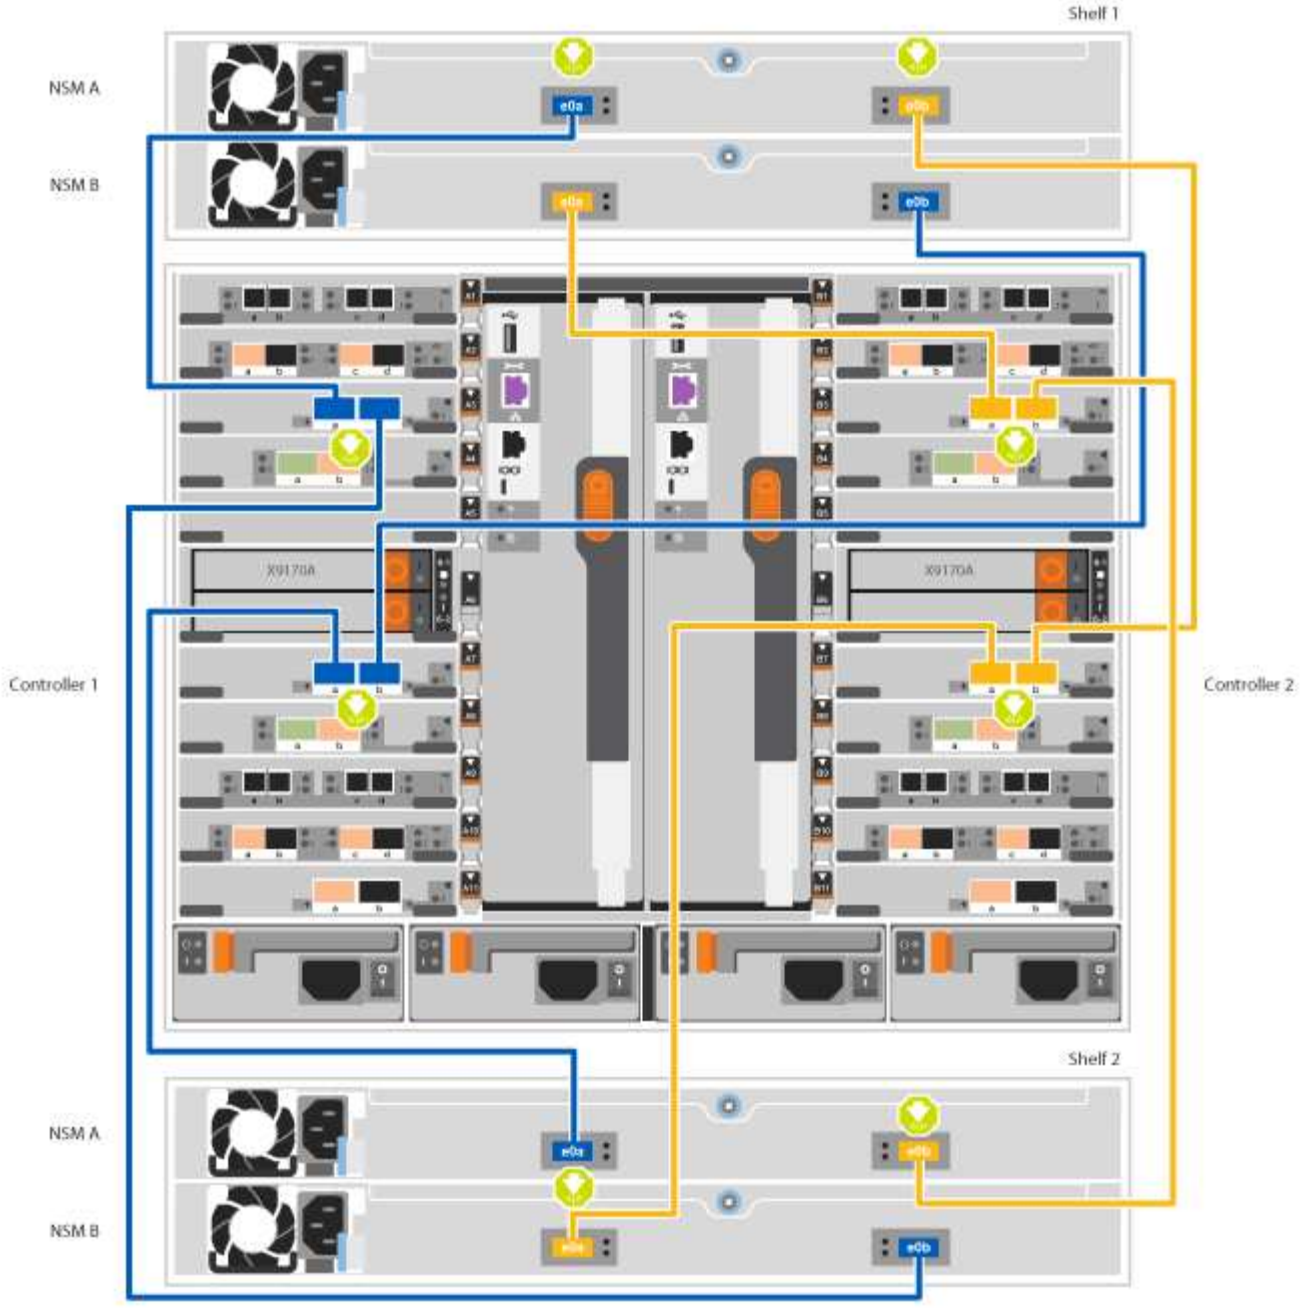

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Completing system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Use the following animation to set one or more drive shelf IDs:

If your system has NS224 drive shelves, the shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located.

Animation - Set SAS or NVMe drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Turn on the power switches to both nodes.

   Animation - Turn on the power to the controllers

   Initial booting may take up to eight minutes.

4. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

5. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

6. Select an ONTAP icon listed to discover:

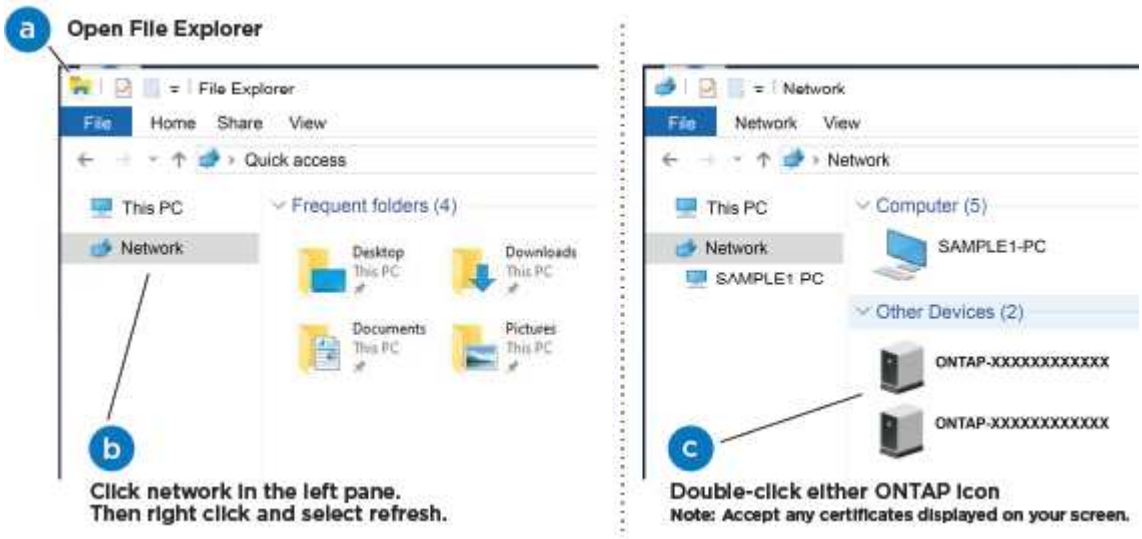

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

   System Manager opens.

7. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

   ONTAP Configuration Guide

8. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

9. Verify the health of your system by running Config Advisor.
10. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet .

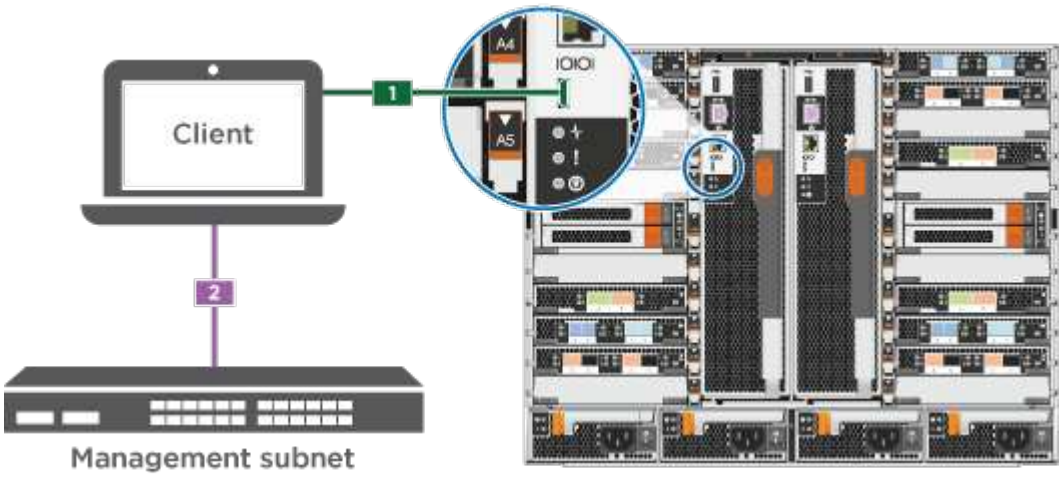

   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Use the following animation to set one or more drive shelf IDs:

   If your system has NS224 drive shelves, the shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located.

   Animation - Set SAS or NVMe drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
4. Turn on the power switches to both nodes.

   Animation - Turn on the power to the controllers

   Initial booting may take up to eight minutes.

5. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

6. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the *NetApp ONTAP Configuration guide*.

      ONTAP Configuration Guide

7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain AFF A700 hardware

For the AFF A700 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Caching module

You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DCPM

The DCPM (destage controller power module) contains the NVRAM11 battery.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Fan

The fan cools the controller.

I/O module

The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.

LED USB

The LED USB module provides connectivity to console ports and system status.

NVRAM

The NVRAM module (Non-Volatile Random Access Memory) allows the controller to retain data across power cycles or system reboots.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real-time clock battery

A real time clock battery preserves system date and time information if the power is off.

X91148A module

The X91148A module is an I/O module that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.

Boot media

Overview of boot media replacement - AFF A700 and FAS9000

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz`.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair does not require connection to a network to restore the `var` file system. The HA pair in a single chassis has an internal e0S connection, which is used to transfer `var` config between them.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct node:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy node* is the HA partner of the impaired node.

Check encryption key support and status - AFF A700

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the impaired controller - AFF A700

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.

If the impaired controller displays…	Then…
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired node.
NOTE: Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 3: Controller is in a two-node MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired node.

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - AFF A700

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the orange button on the cam handle downward until it unlocks.

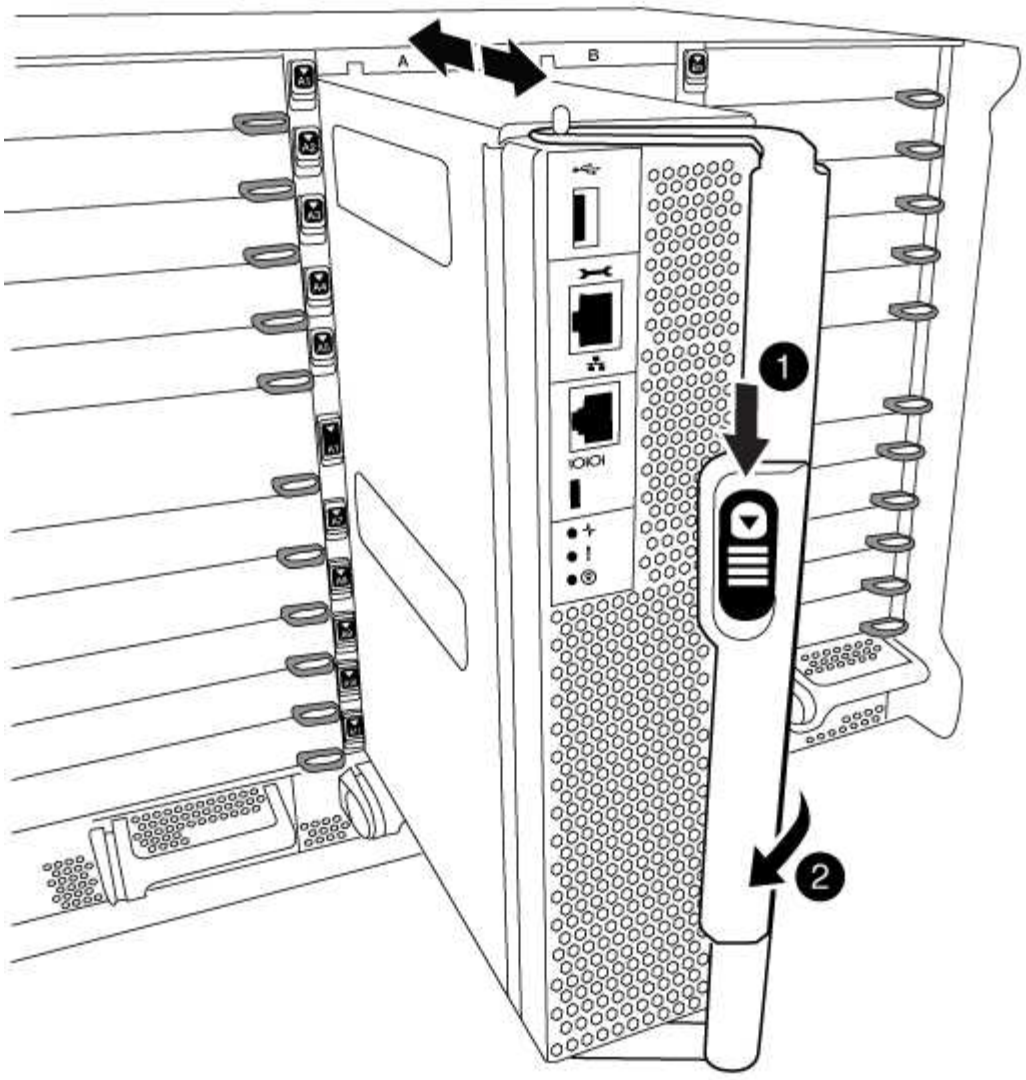

1	
Cam handle release button	
2	
Cam handle	

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

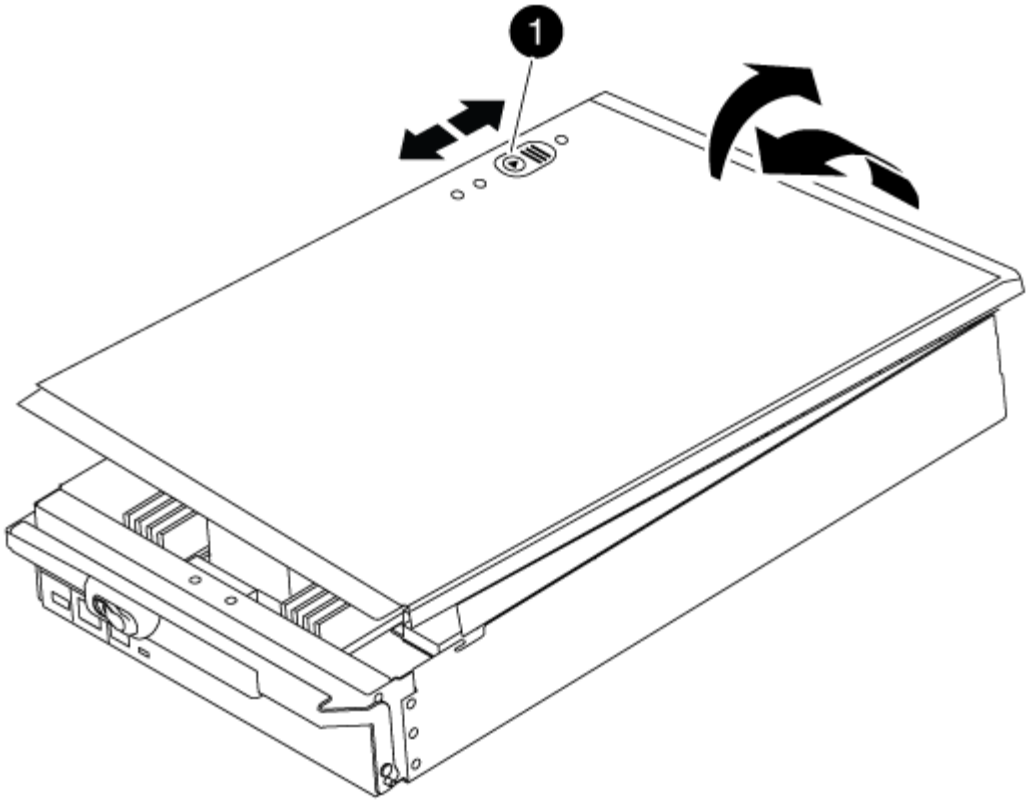

1	Controller module cover locking button

Step 2: Replace the boot media

Locate the boot media using the following illustration or the FRU map on the controller module:

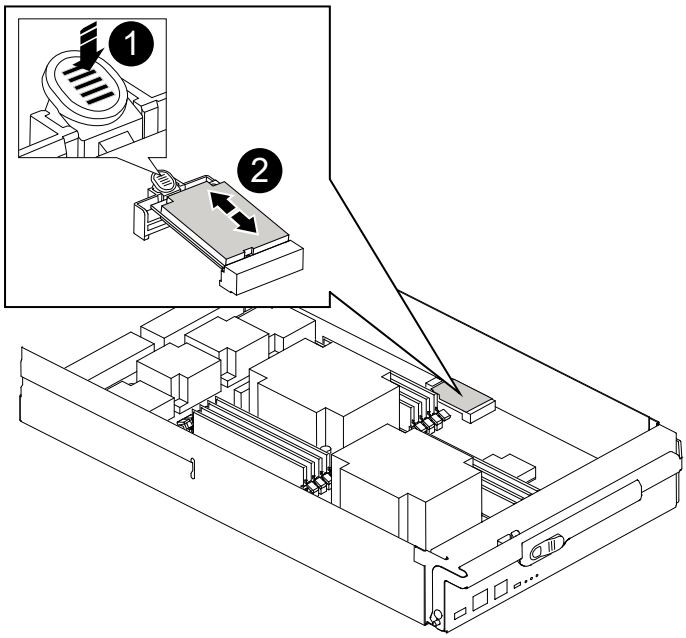

1	
Press release tab	
2	
Boot media	

1. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

2. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
3. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

4. Push the boot media down to engage the locking button on the boot media housing.
5. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the `var` file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the `var` file system.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Recable the controller module, as needed.
3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB

console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, and then push the cam handle to the closed position.

   The node begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the node to boot to LOADER.

6. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

     The target port you configure is the target port you use to communicate with the impaired node from the healthy node during `var` file system restore with a network connection. You can also use the e0M port in this command.

   - If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
     - filer_addr is the IP address of the storage system.
     - netmask is the network mask of the management network that is connected to the HA partner.
     - gateway is the gateway for the network.
     - dns_addr is the IP address of a name server on your network.
     - dns_domain is the Domain Name System (DNS) domain name.

       If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

     Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

7. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:

   a. Boot to Maintenance mode: `boot_ontap maint`

   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`

   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Boot the recovery image - AFF A700

The procedure for booting the impaired node from the recovery image depends on whether the system is in a two-node MetroCluster configuration.

Option 1: Boot the recovery image in most systems

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

This procedure applies to systems that are not in a two-node MetroCluster configuration.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the `var` file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy node to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the node to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the node.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

If your system has…	Then…

No network connection and is in a MetroCluster IP configuration	a. Press `n` when prompted to restore the backup configuration.
	b. Reboot the system when prompted by the system.
	c. Wait for the iSCSI storage connections to connect.
	You can proceed after you see the following messages:
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_auxiliary, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_partner, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_auxiliary, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_partner, address: ip-address).`
	d. Select the **Update flash from backup config** (sync flash) option from the displayed menu.
	If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the node to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:
   - If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed

- If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

*If you see…	Then…*
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner node. b. Confirm the target node is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner node.
8. Give back the node using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired node and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Option 2: Boot the recovery image in a two-node MetroCluster configuration

You must boot the ONTAP image from the USB drive and verify the environmental variables.

This procedure applies to systems in a two-node MetroCluster configuration.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. After the image is installed, start the restoration process:
   a. Press `n` when prompted to restore the backup configuration.
   b. Press `y` when prompted to reboot to start using the newly installed software.

      You should be prepared to interrupt the boot process when prompted.

4. As the system boots, press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message., and when the Boot Menu is displayed select option 6.
5. Verify that the environmental variables are set as expected.
   a. Take the node to the LOADER prompt.

b. Check the environment variable settings with the `printenv` command.

c. If an environment variable is not set as expected, modify it with the `setenv *environment-variable-name changed-value*` command.

d. Save your changes using the `savenv` command.

e. Reboot the node.

Switch back aggregates in a two-node MetroCluster configuration - AFF A700 and FAS9000

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Restore encryption - AFF A700

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF A700 and FAS9000

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF A700

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A700

To replace the chassis, you must shutdown the controllers.

Option 1: Shut down the controllers

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true
   ```

For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Shut down a node in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

   If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Move the fans, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Remove the power supplies

Steps

Removing the power supplies when replacing a chassis involves turning off, disconnecting, and then removing the power supply from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Press and hold the orange button on the power supply handle, and then pull the power supply out of the chassis.

   When removing a power supply, always use two hands to support its weight.

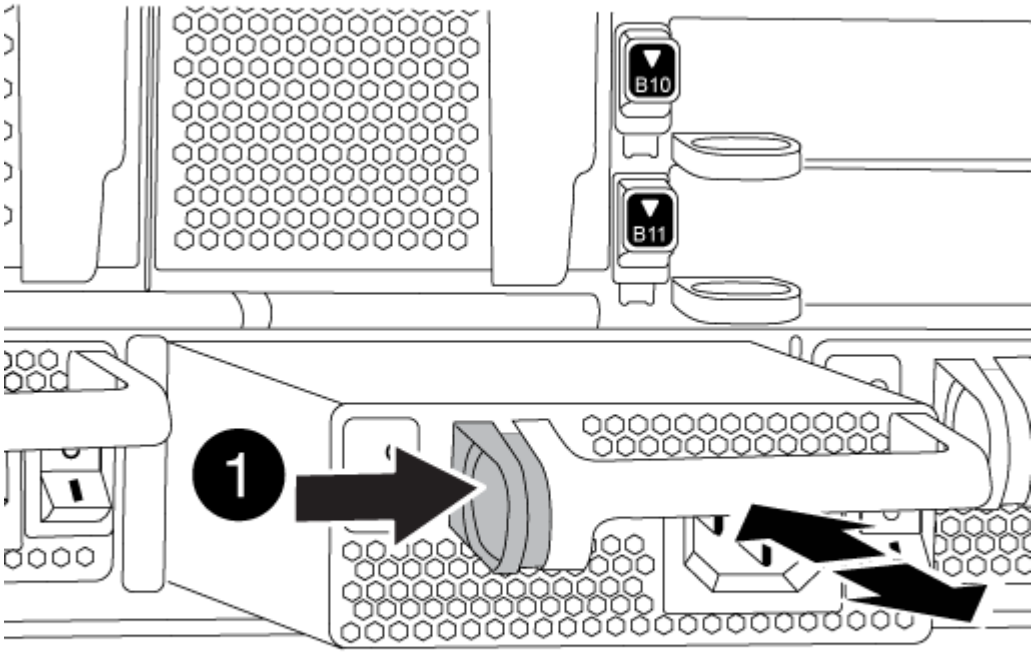

1	Locking button

4. Repeat the preceding steps for any remaining power supplies.

Step 2: Remove the fans

To remove the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

Steps

1. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
2. Press the orange button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

> The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

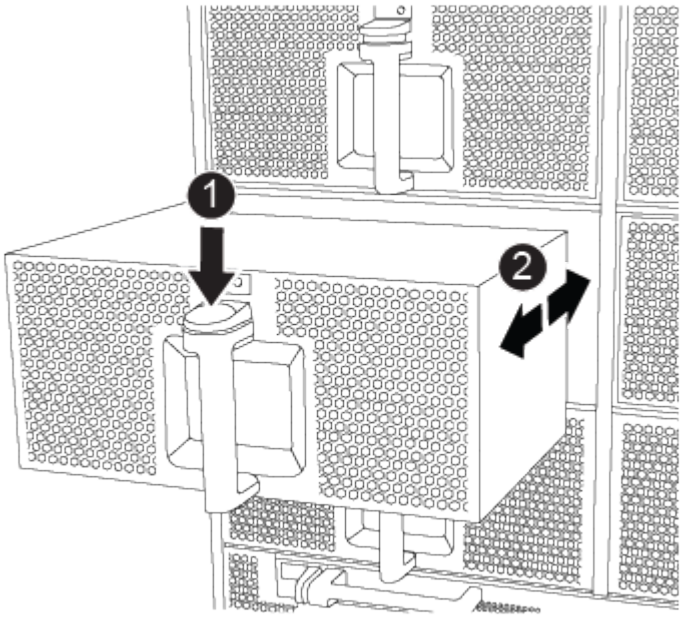

1	Orange release button

3. Set the fan module aside.
4. Repeat the preceding steps for any remaining fan modules.

Step 3: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the old chassis.

Steps

1. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.

2. Slide the orange button on the cam handle downward until it unlocks.

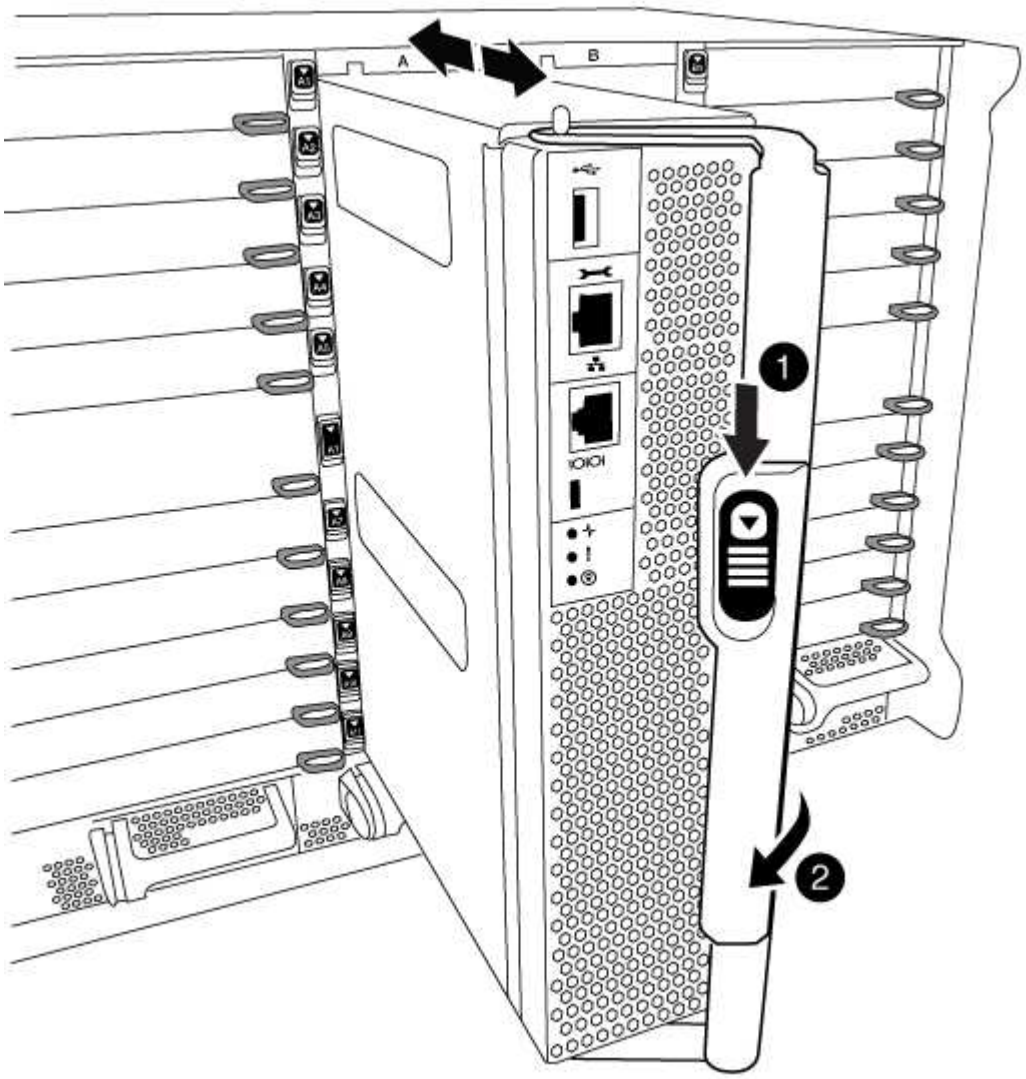

1	Cam handle release button
2	Cam handle

3. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

4. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 4: Remove the I/O modules

Steps

To remove I/O modules from the old chassis, including the NVRAM modules, follow the specific sequence of

steps. You do not have to remove the Flash Cache module from the NVRAM module when moving it to a new chassis.

1. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

2. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

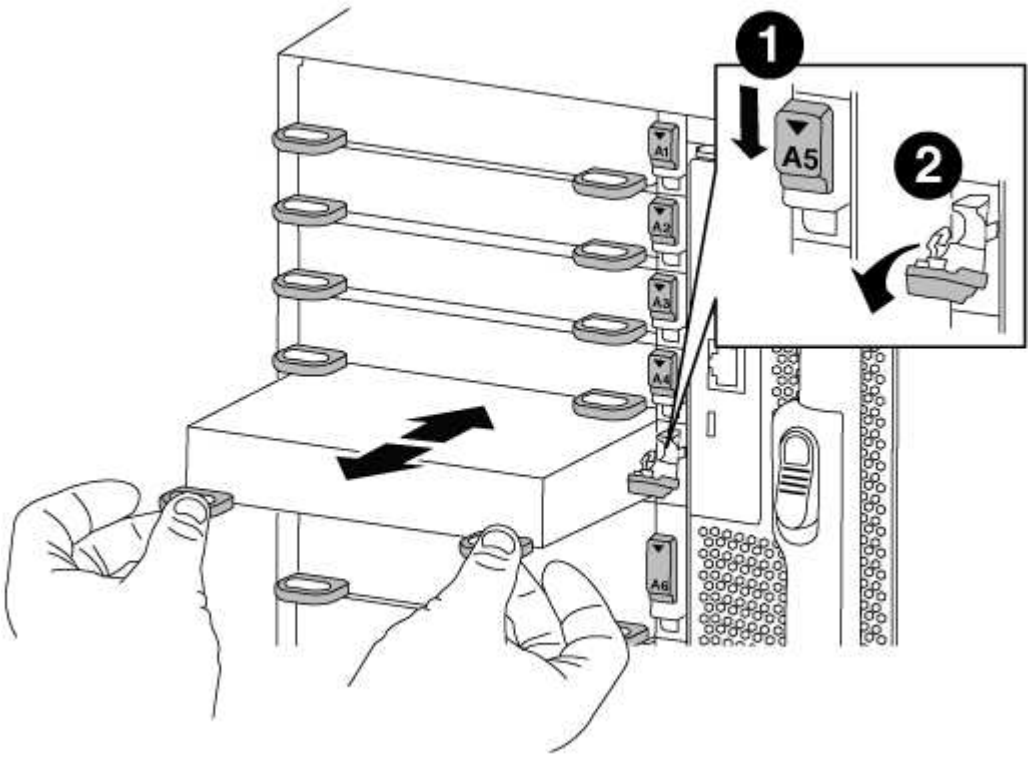

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

3. Set the I/O module aside.
4. Repeat the preceding step for the remaining I/O modules in the old chassis.

Step 5: Remove the De-stage Controller Power Module

Steps

You must remove the de-stage controller power modules from the old chassis in preparation for installing the replacement chassis.

1. Press the orange locking button on the module handle, and then slide the DCPM module out of the chassis.

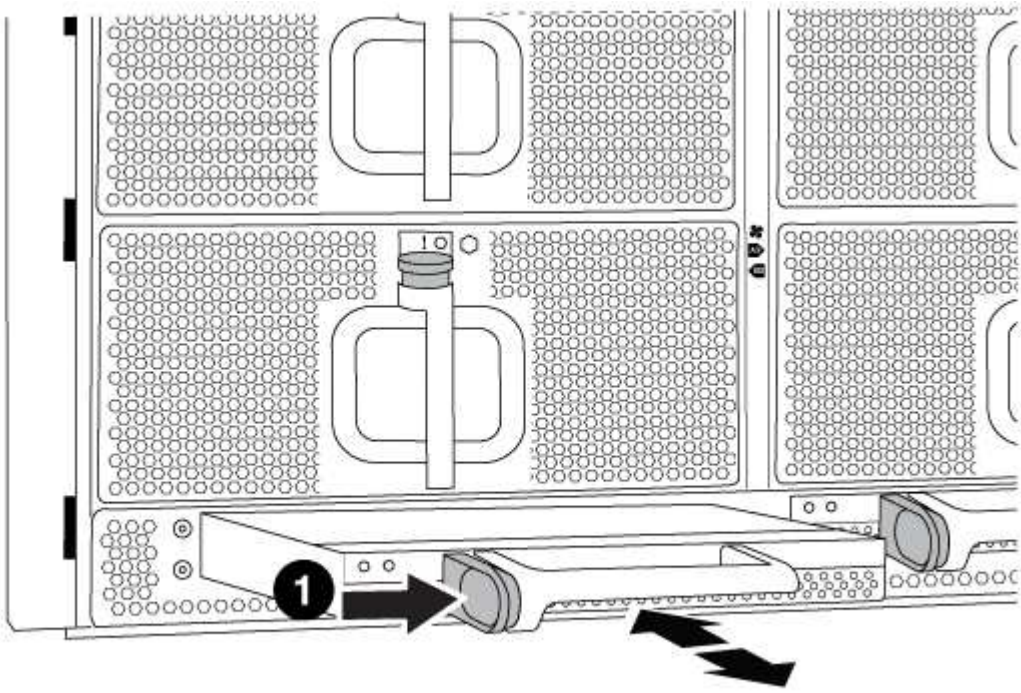

1	DCPM module orange locking button

2. Set the DCPM module aside in a safe place and repeat this step for the remaining DCPM module.

Step 6: Replace a chassis from within the equipment rack or system cabinet

Steps

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.

   If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.

5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. Secure the rear of the chassis to the equipment rack or system cabinet.
8. If you are using the cable management brackets, remove them from the old chassis, and then install them on the replacement chassis.
9. If you have not already done so, install the bezel.

Step 7: Move the USB LED module to the new chassis

Steps

Once the new chassis is installed into the rack or cabinet, you must move the USB LED module from the old chassis to the new chassis.

1. Locate the USB LED module on the front of the old chassis, directly under the power supply bays.
2. Press the black locking button on the right side of the module to release the module from the chassis, and then slide it out of the old chassis.
3. Align the edges of the module with the USB LED bay at the bottom-front of the replacement chassis, and gently push the module all the way into the chassis until it clicks into place.

Step 8: Install the de-stage controller power module when replacing the chassis

Steps

Once the replacement chassis is installed into the rack or system cabinet, you must reinstall the de-stage controller power modules into it.

1. Align the end of the DCPM module with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

2. Repeat this step for the remaining DCPM module.

Step 9: Install fans into the chassis

Steps

To install the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

1. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

2. Repeat these steps for the remaining fan modules.
3. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Step 10: Install I/O modules

Steps

To install I/O modules, including the NVRAM/Flash Cache modules from the old chassis, follow the specific sequence of steps.

You must have the chassis installed so that you can install the I/O modules into the corresponding slots in the new chassis.

1. After the replacement chassis is installed in the rack or cabinet, install the I/O modules into their corresponding slots in the replacement chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage, and then push the I/O cam latch all the way up to lock the module in place.
2. Recable the I/O module, as needed.
3. Repeat the preceding step for the remaining I/O modules that you set aside.

   If the old chassis has blank I/O panels, move them to the replacement chassis at this time.

Step 11: Install the power supplies

Steps

Installing the power supplies when replacing a chassis involves installing the power supplies into the replacement chassis, and connecting to the power source.

1. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

2. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

   Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

3. Repeat the preceding steps for any remaining power supplies.

Step 12: Install the controller

Steps

After you install the controller module and any other components into the new chassis, boot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.

3. Connect the power supplies to different power sources, and then turn them on.
4. With the cam handle in the open position, slide the controller module into the chassis and firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle until it clicks into the locked position.

   > Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.

   The controller module begins to boot as soon as it is fully seated in the chassis.

5. Repeat the preceding steps to install the second controller into the new chassis.
6. Boot each node to Maintenance mode:
   a. As each node starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

      > If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

   b. From the boot menu, select the option for Maintenance mode.

Complete the restoration and replacement process - AFF A700

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for `HA-state` can be one of the following:

      - `ha`
      - `mcc`
      - `mcc-2n`
      - `mccip`
      - `non-ha`

b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.
4. Exit Maintenance mode: `halt`

   The LOADER prompt appears.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller module

Overview of controller module replacement - AFF A700

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system has a V_StorageAttach license, you must refer to the additional required steps before performing this procedure.
- If your system is in an HA pair, the healthy node must be able to take over the node that is being replaced (referred to in this procedure as the “impaired node”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a node in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired node to the *replacement* node so that the *replacement* node will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* node is the node that is being replaced.
  - The *replacement* node is the new node that is replacing the impaired node.
  - The *healthy* node is the surviving node.

- You must always capture the node's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - AFF 700

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the controller module hardware - AFF A700

To replace the controller module hardware, you must remove the impaired node, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the orange button on the cam handle downward until it unlocks.

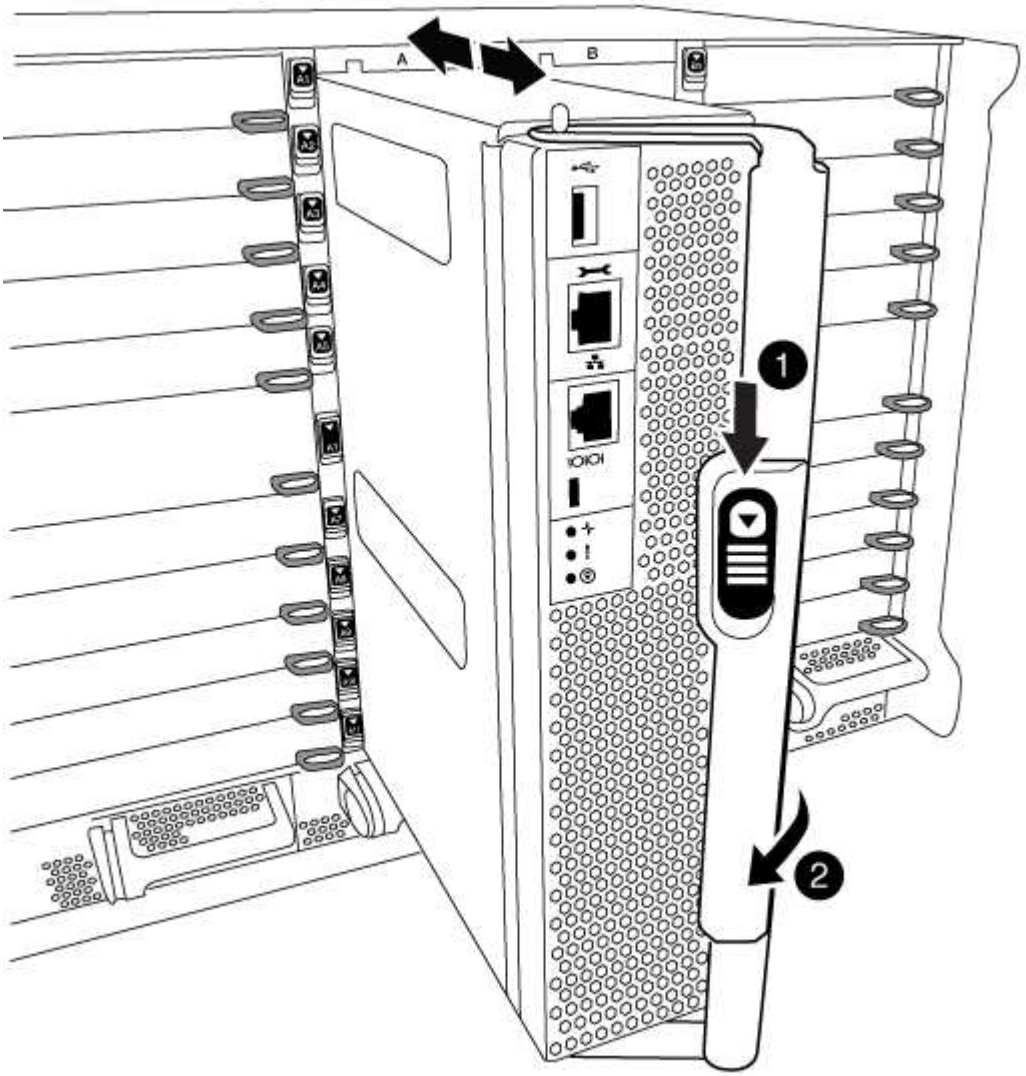

1	Cam handle release button
2	Cam handle

1. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

2. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

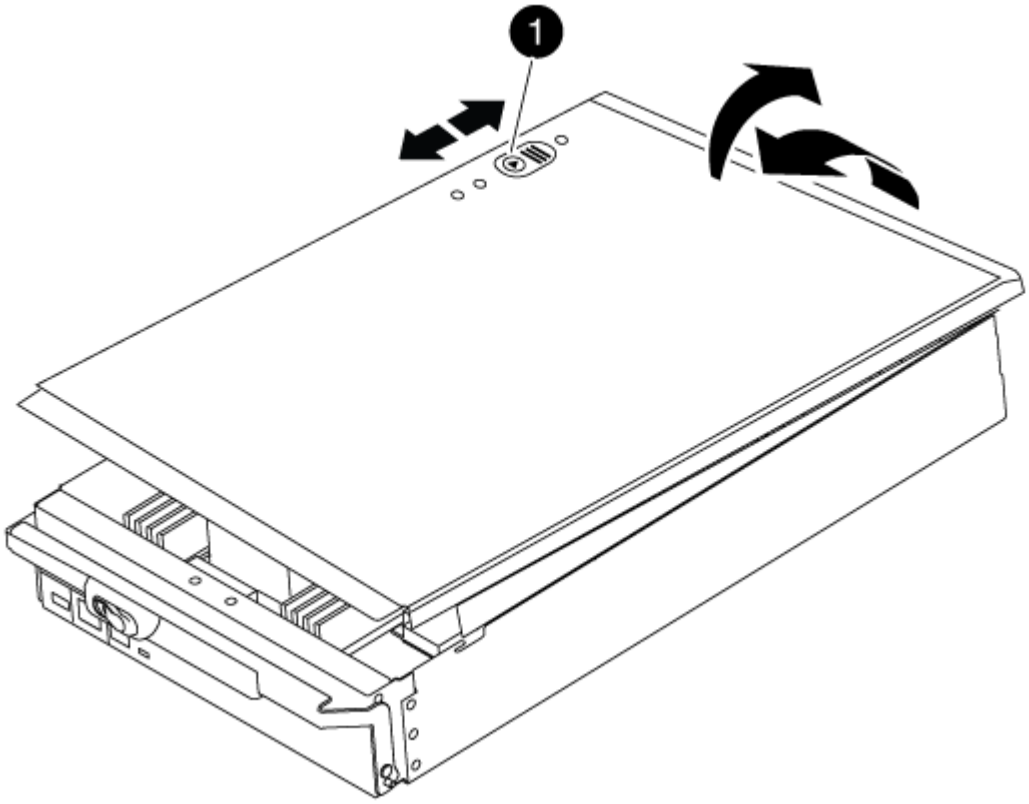

1	Controller module cover locking button

Step 2: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller and insert it in the new controller.

Steps

1. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

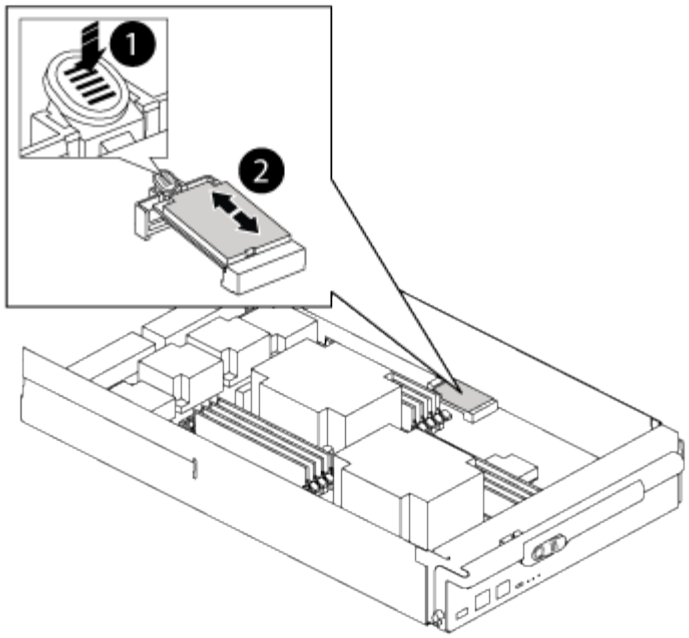

1	Press release tab
2	Boot media

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the system DIMMs

To move the DIMMs, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM,

and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

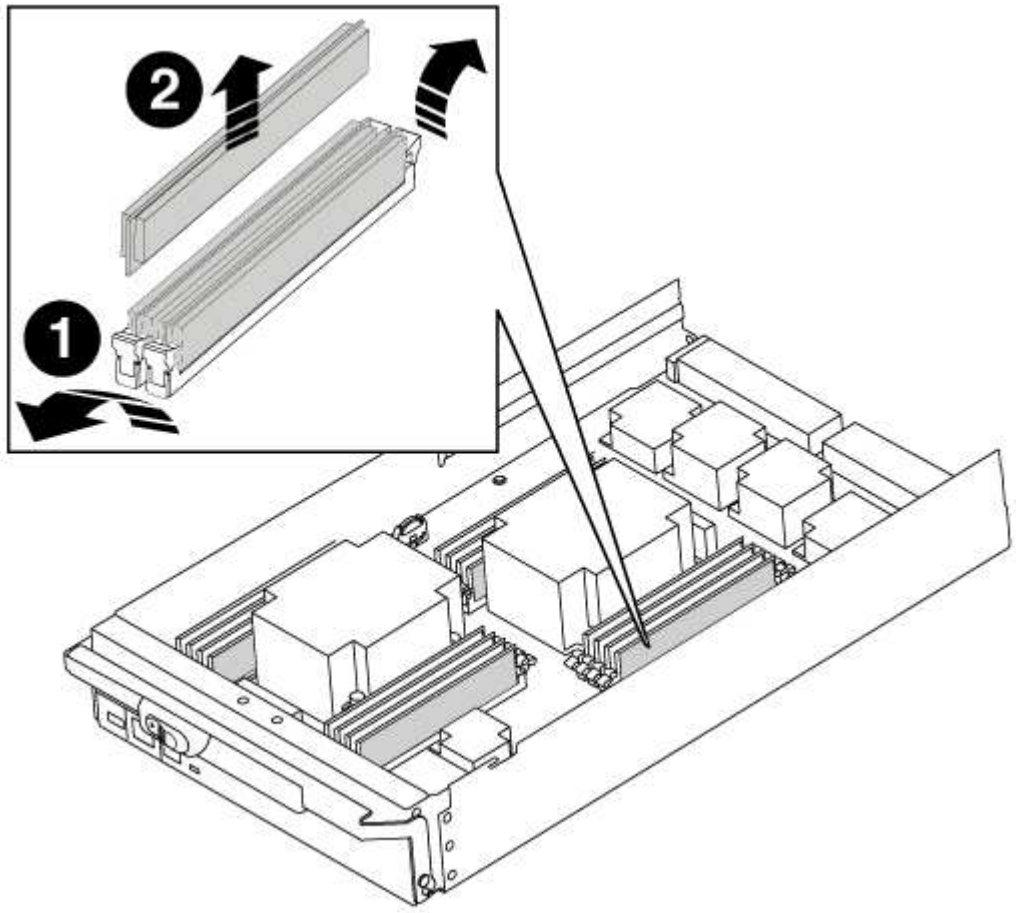

1	DIMM ejector tabs
2	DIMM

5. Locate the slot where you are installing the DIMM.
6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
9. Repeat these steps for the remaining DIMMs.

Step 4: Install the controller

After you install the components into the controller module, you must install the controller module back into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   d. Interrupt the boot process by pressing `Ctrl-C` when you see `Press Ctrl-C for Boot Menu`.
   e. Select the option to boot to Maintenance mode from the displayed menu.

Restore and verify the system configuration - AFF A700

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the controller's HA state

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`

- `mccip`
- `non-ha`

  a. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF A700

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* node and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* node is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* node, boot the node, entering `y` if you are prompted to override the system ID due to a system ID mismatch.`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* node console and then, from the healthy node, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired node, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node                Partner         Possible     State Description
------------      ------------      --------
-------------------------------------
node1               node2           false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2               node1           -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy node, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the node:
   a. From the healthy node, give back the replaced node's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* node takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration Guide for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed

on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* node should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the node: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each node will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the node is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a node on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The *replacement* node is the current owner of the disks on the disaster site.

     Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each node is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each node: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy node: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF A700

To complete the replacement procedure and restore your system to full operation, you must recable the storage, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Install licenses for the replacement node in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

  Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The licenses keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.
- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verifying LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: (MetroCluster only): Switching back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a de-stage controller power module (DCPM) - AFF A700

To hot-swap a de-stage controller power module (DCPM), which contains the NVRAM10 battery, you must locate the failed DCPM module, remove it from the chassis, and install the replacement DCPM module.

You must have a replacement DCPM module in-hand before removing the failed module from the chassis and it must be replaced within five minutes of removal. Once the DCPM module is removed from the chassis, there is no shutdown protection for the controller module that owns the DCPM module, other than failover to the other controller module.

Step 1: Replace the DCPM module

To replace the DCPM module in your system, you must remove the failed DCPM module from the system and then replace it with a new DCPM module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel on the front of the system and set it aside.
3. Locate the failed DCPM module in the front of the system by looking for the Attention LED on the module.

   The LED will be steady amber if the module is faulty.

   > The DCPM module must be replaced in the chassis within five minutes of removal or the associated controller will shut down.

4. Press the orange locking button on the module handle, and then slide the DCPM module out of the chassis.

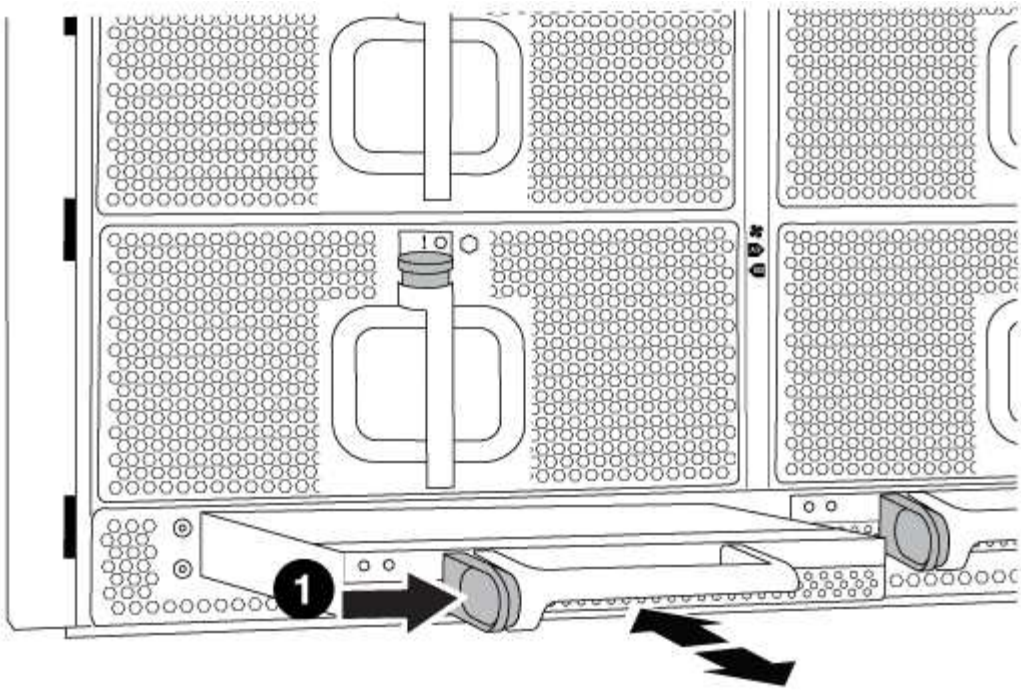

1	DCPM module orange locking button

5. Align the end of the DCPM module with the chassis opening, and then gently slide it into the chassis until it clicks into place.

> The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

The DCPM module LED lights when the module is fully seated into the chassis.

Step 2: Dispose of batteries

You must dispose of batteries according to the local regulations regarding battery recycling or disposal. If you cannot properly dispose of batteries, you must return the batteries to NetApp, as described in the RMA instructions that are shipped with the kit.

https://library.netapp.com/ecm/ecm_download_file/ECMP12475945

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A700

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM

failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the orange button on the cam handle downward until it unlocks.

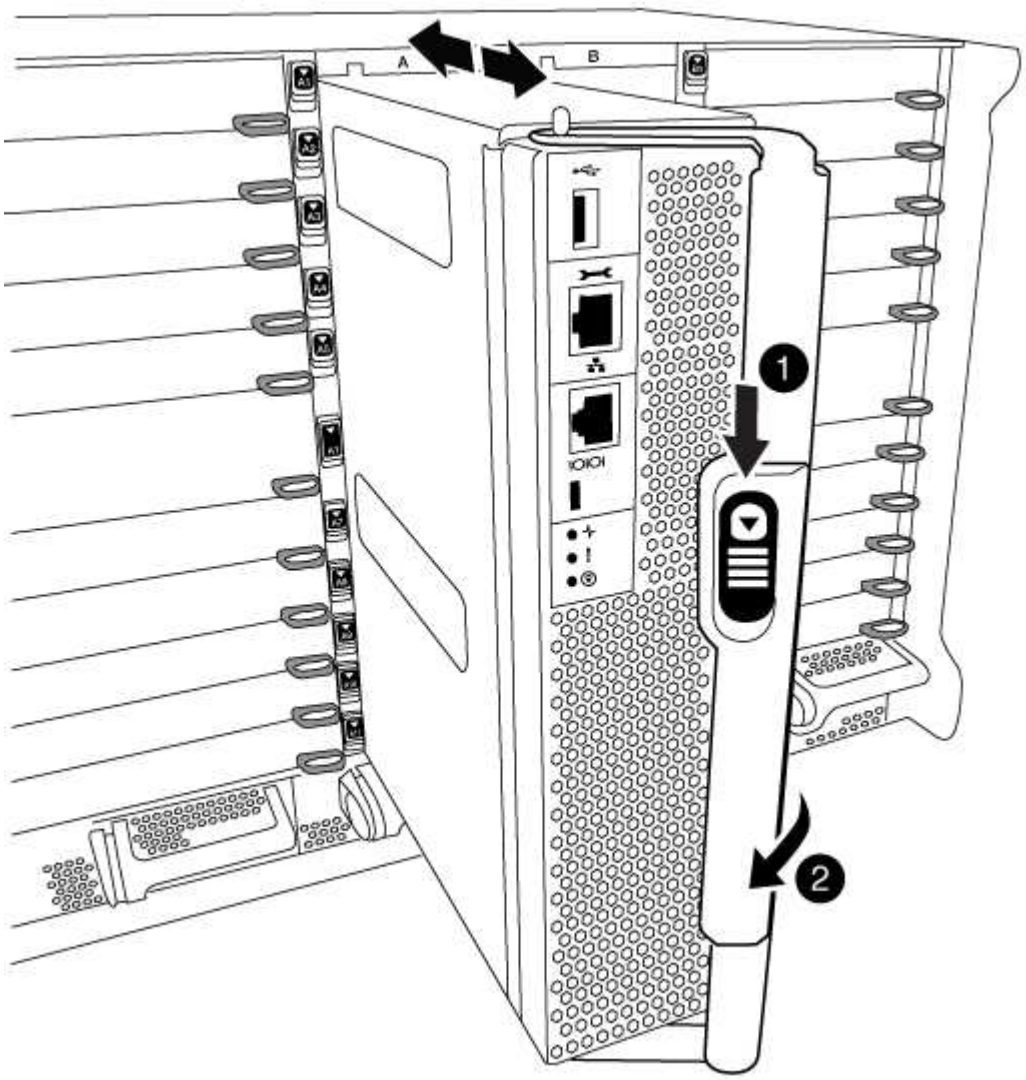

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

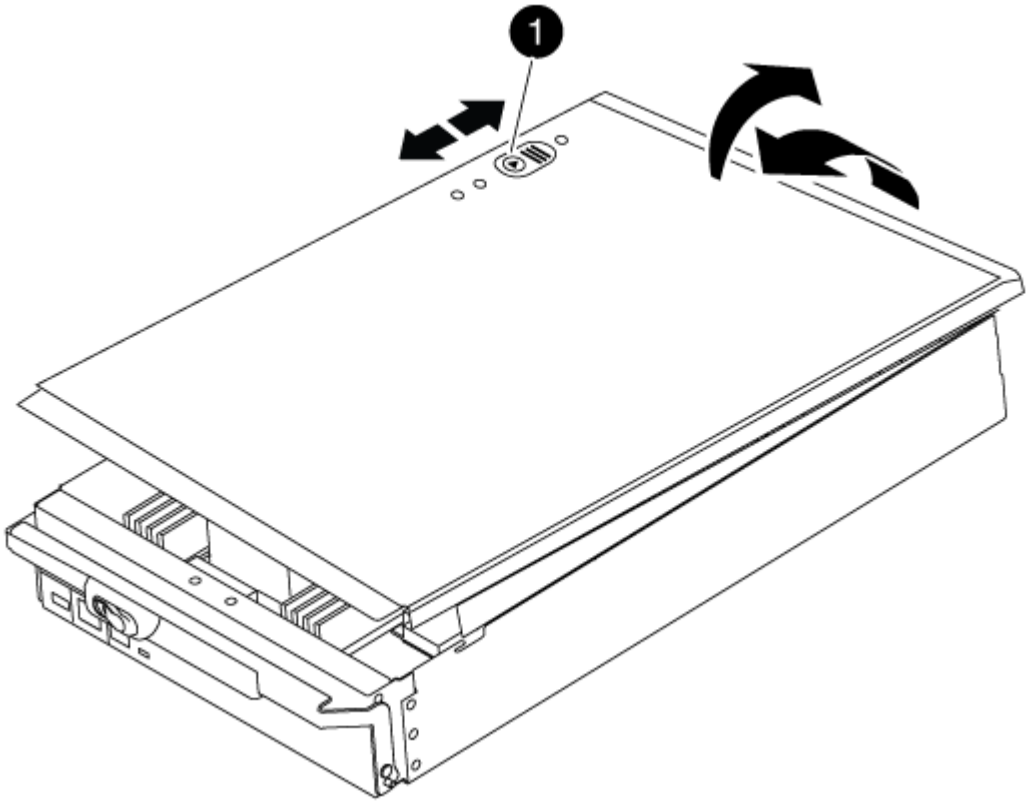

1	Controller module cover locking button

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.

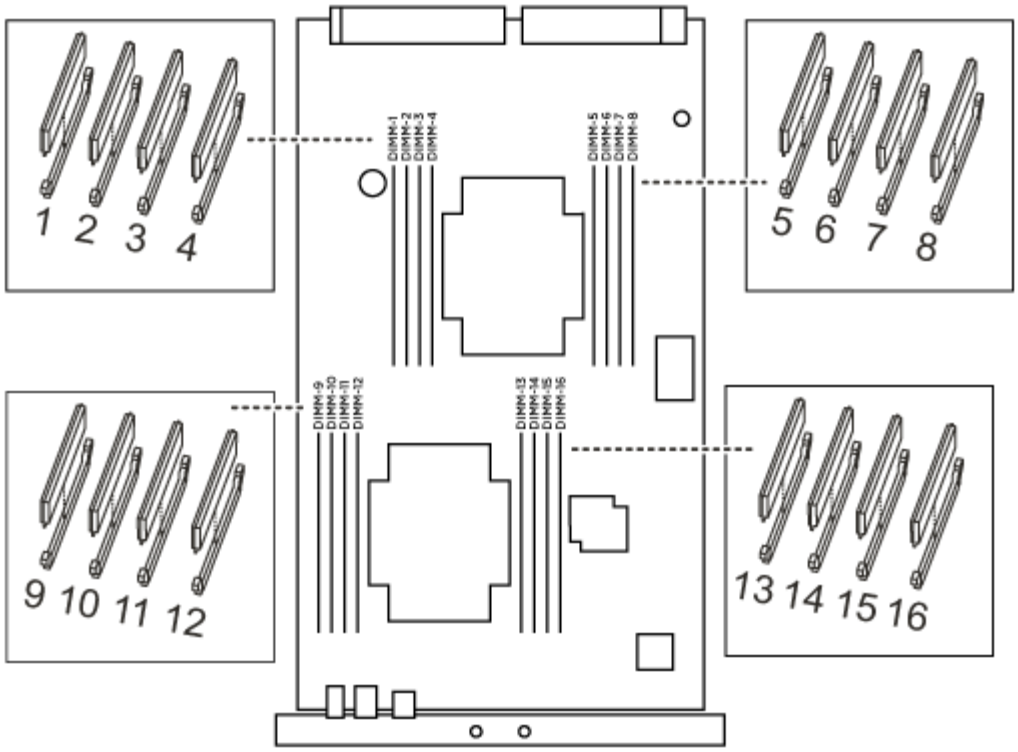

1. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

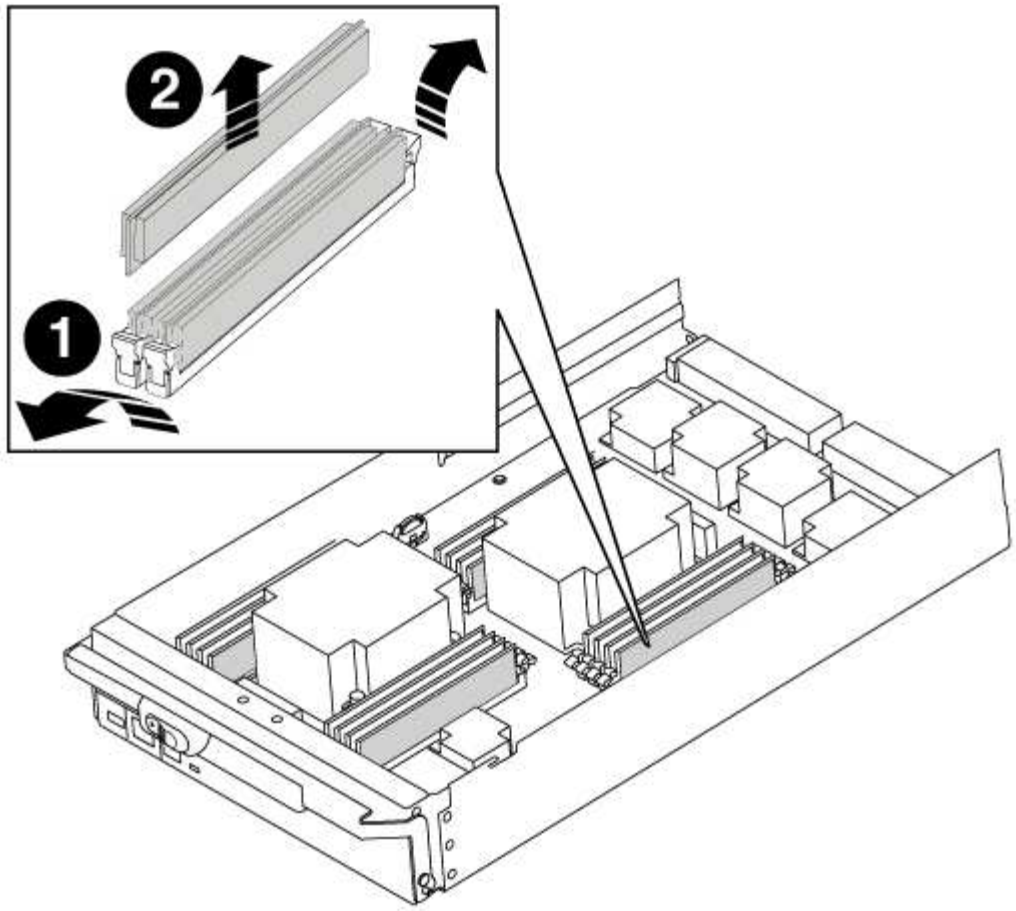

1	DIMM ejector tabs
2	DIMM

2. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

3. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

4. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
5. Close the controller module cover.

Step 4: Install the controller

After you install the components into the controller module, you must install the controller module back into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a fan - AFF A700

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press the orange button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

> The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

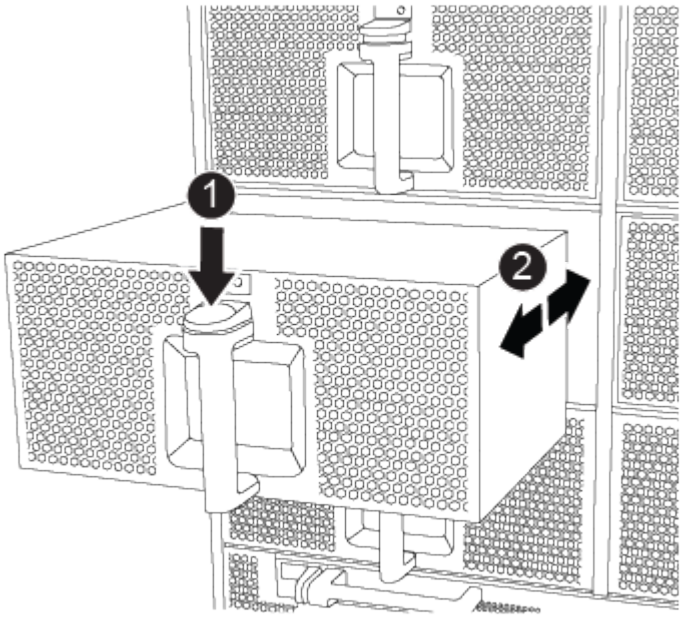

1	Orange release button

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the

chassis until it snaps into place.

When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an I/O module - AFF A700 and FAS9000

To replace an I/O module, you must perform a specific sequence of tasks.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Replace I/O modules

To replace an I/O module, locate it within the chassis and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

3. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

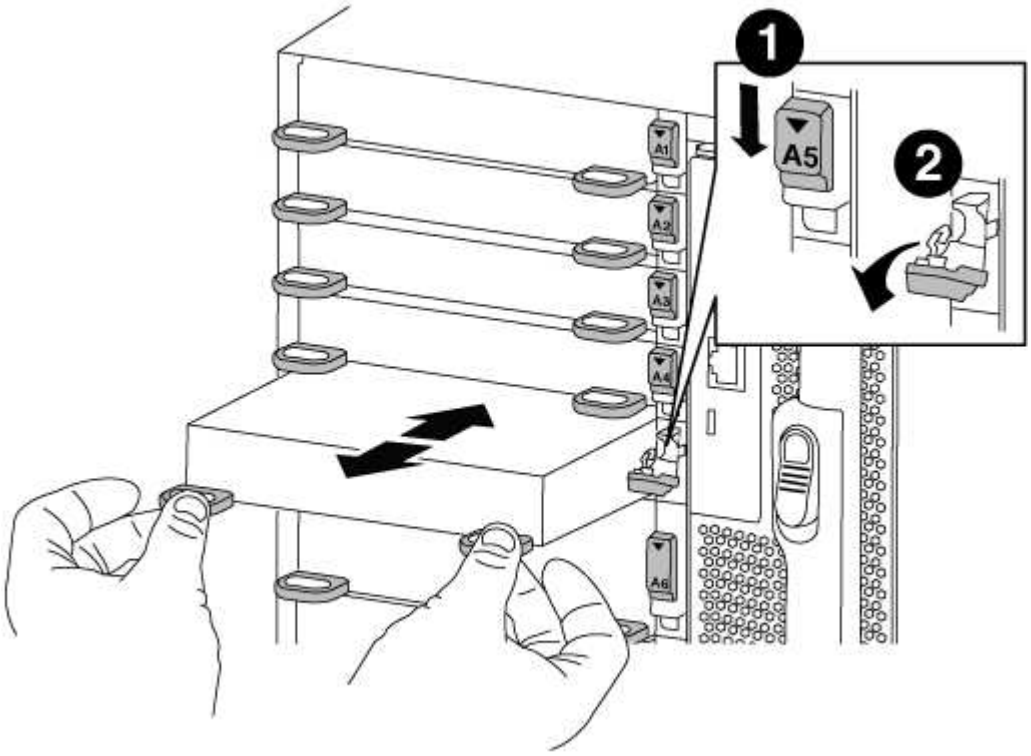

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Install the replacement I/O module into the chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.
6. Recable the I/O module, as needed.

Step 3: Reboot the controller after I/O module replacement

After you replace an I/O module, you must reboot the controller module.

> If the new I/O module is not the same model as the failed module, you must first reboot the BMC.

Steps

1. Reboot the BMC if the replacement module is not the same model as the old module:
   a. From the LOADER prompt, change to advanced privilege mode: `priv set advanced`
   b. Reboot the BMC: `sp reboot`
2. From the LOADER prompt, reboot the node: `bye`

   > This reinitializes the PCIe cards and other components and reboots the node.
3. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs or onboard ports, convert these ports to 10 GbE connections by using the `nicadmin convert` command from Maintenance mode.

   > Be sure to exit Maintenance mode after completing the conversion.
4. Return the node to normal operation:
   `storage failover giveback -ofnode` *`impaired_node_name`*
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

   > If your system is in a two-node MetroCluster configuration, you must switch back the aggregates as described in the next step.

Step 4: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an LED USB module - AFF A700

You can replace an LED USB module without interrupting service.

The FAS9000 or AFF A700 LED USB module provides connectivity to console ports and system status. Replacement of this module does not require tools.

Steps

1. Remove the old LED USB module:

   a. With the bezel removed, locate the LED USB module at the front of the chassis, on the bottom left side.
   b. Slide the latch to partially eject the module.
   c. Pull the module out of the bay to disconnect it from the midplane. Do not leave the slot empty.

2. Install the new LED USB module:

   a. Align the module to the bay with the notch in the corner of the module positioned near the slider latch on the chassis. The bay will prevent you from installing the module upside down.

b. Push the module into the bay until it is fully seated flush with the chassis.

   There is an audible click when the module is secure and connected to the midplane.

Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVRAM module or NVRAM DIMMs - AFF A700

The NVRAM module consists of the NVRAM10 and DIMMs and up to two NVMe SSD Flash Cache modules (Flash Cache or caching modules) per NVRAM module. You can replace a failed NVRAM module or the DIMMs inside the NVRAM module.

To replace a failed NVRAM module, you must remove it from the chassis, remove the Flash Cache module or modules from the NVRAM module, move the DIMMs to the replacement module, reinstall the Flash Cache module or modules, and install the replacement NVRAM module into the chassis.

Because the system ID is derived from the NVRAM module, if replacing the module, disks belonging to the system are reassigned to the new system ID.

Before you begin

- All disk shelves must be working properly.
- If your system is in an HA pair, the partner node must be able to take over the node associated with the NVRAM module that is being replaced.
- This procedure uses the following terminology:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy* node is the HA partner of the impaired node.
- This procedure includes steps for automatically or manually reassigning disks to the controller module associated with the new NVRAM module. You must reassign the disks when directed to in the procedure. Completing the disk reassignment before giveback can cause issues.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You cannot change any disks or disk shelves as part of this procedure.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a Two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 6 in the chassis and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Move the Flash Cache module from the old NVRAM module to the new NVRAM module:

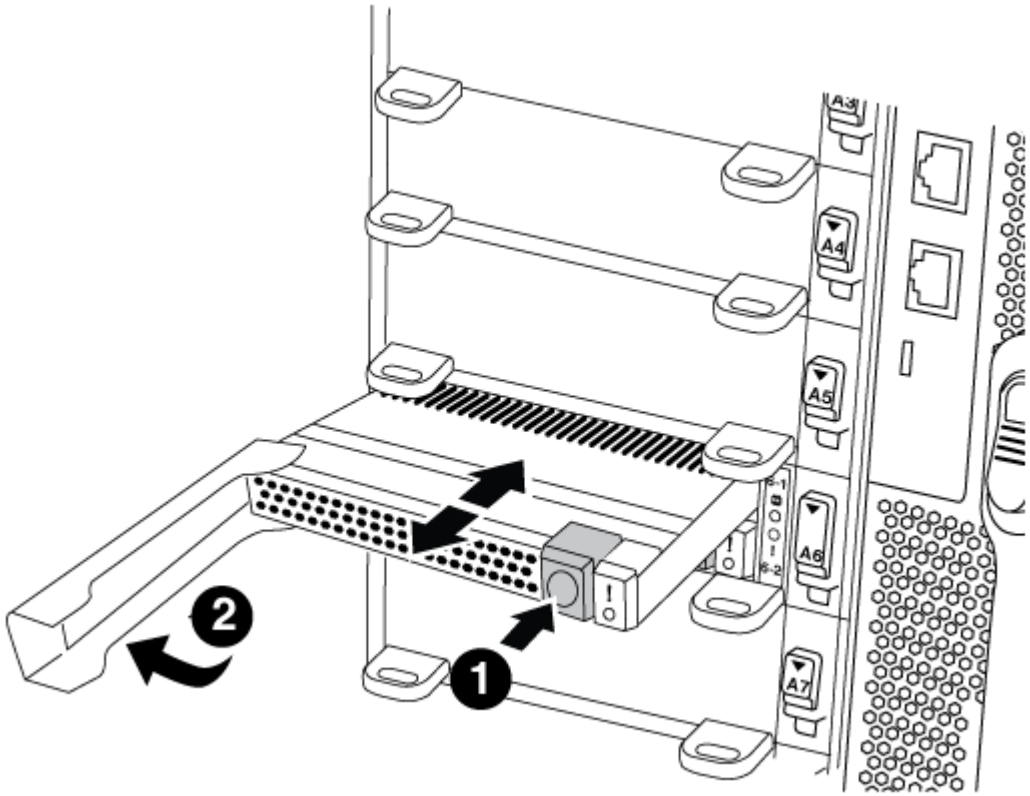

1	Orange release button (gray on empty Flash Cache modules)
2	Flash Cache cam handle

   a. Press the orange button on the front of the Flash Cache module.

      The release button on empty Flash Cache modules is gray.

   b. Swing the cam handle out until the module begins to slide out of the old NVRAM module.

   c. Grasp the module cam handle and slide it out of the NVRAM module and insert it into the front of the new NVRAM module.

   d. Gently push the Flash Cache module all the way into the NVRAM module, and then swing the cam handle closed until it locks the module in place.

3. Remove the target NVRAM module from the chassis:

   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

b. Rotate the cam latch down until it is in a horizontal position.

   The NVRAM module disengages from the chassis and moves out a few inches.

c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

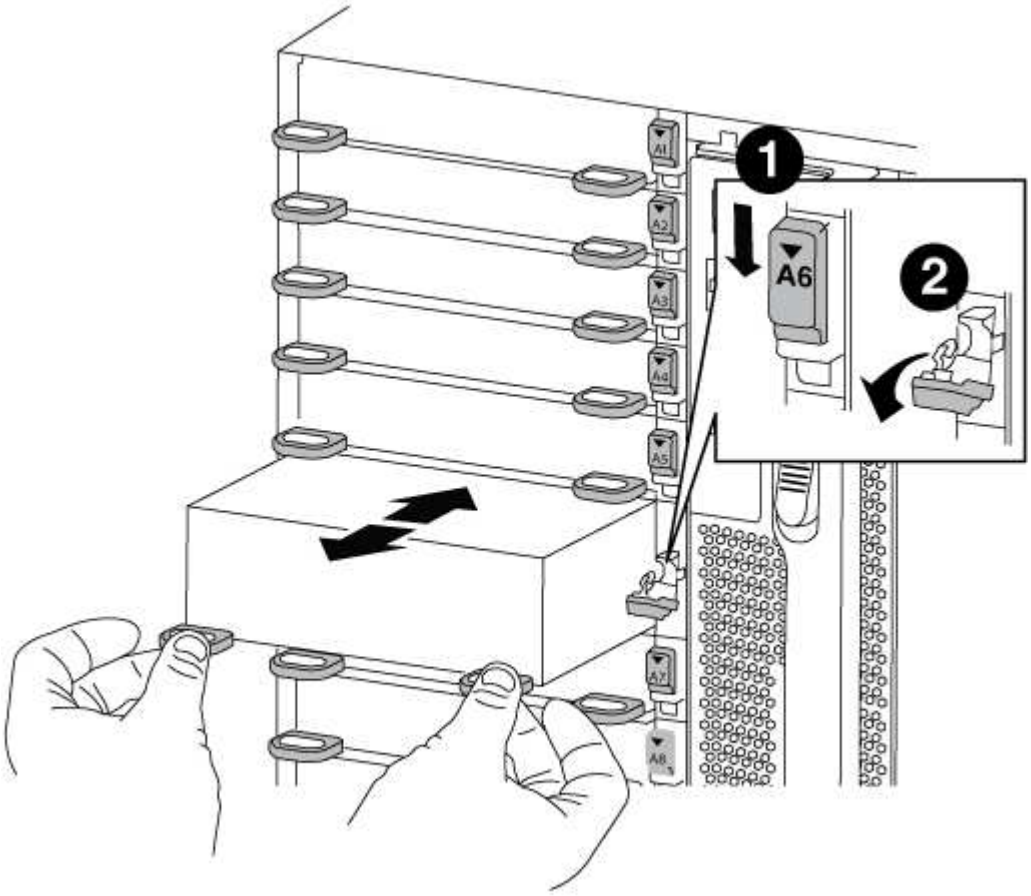

1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

4. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

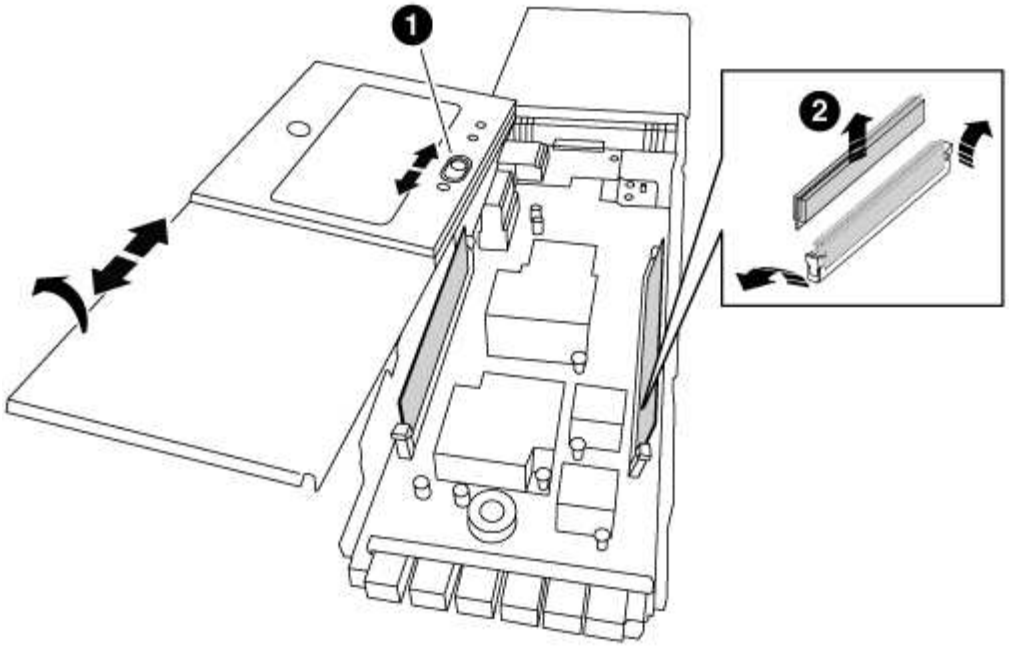

1	Cover locking button
2	DIMM and DIMM ejector tabs

5. Remove the DIMMs, one at a time, from the old NVRAM module and install them in the replacement NVRAM module.
6. Close the cover on the module.
7. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

Step 3: Replace a NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, you must remove the NVRAM module, open the module, and then replace the target DIMM.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the target NVRAM module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The NVRAM module disengages from the chassis and moves out a few inches.

   c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module

face.

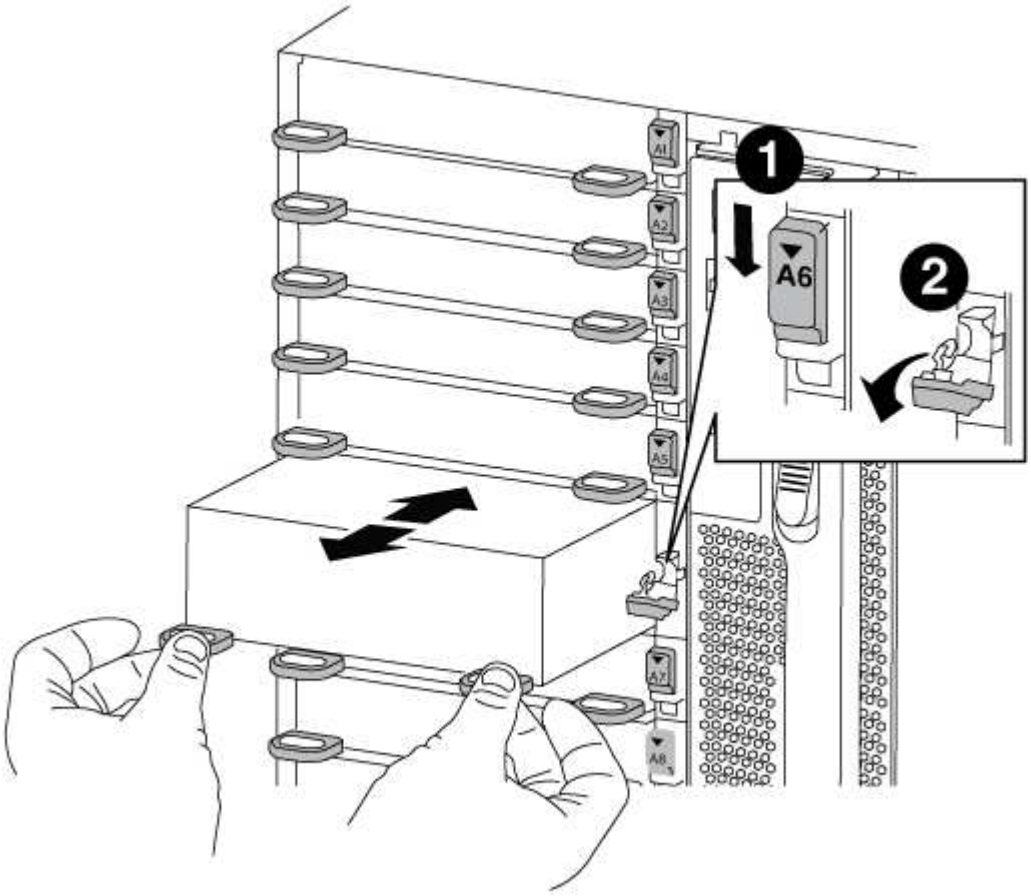

1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

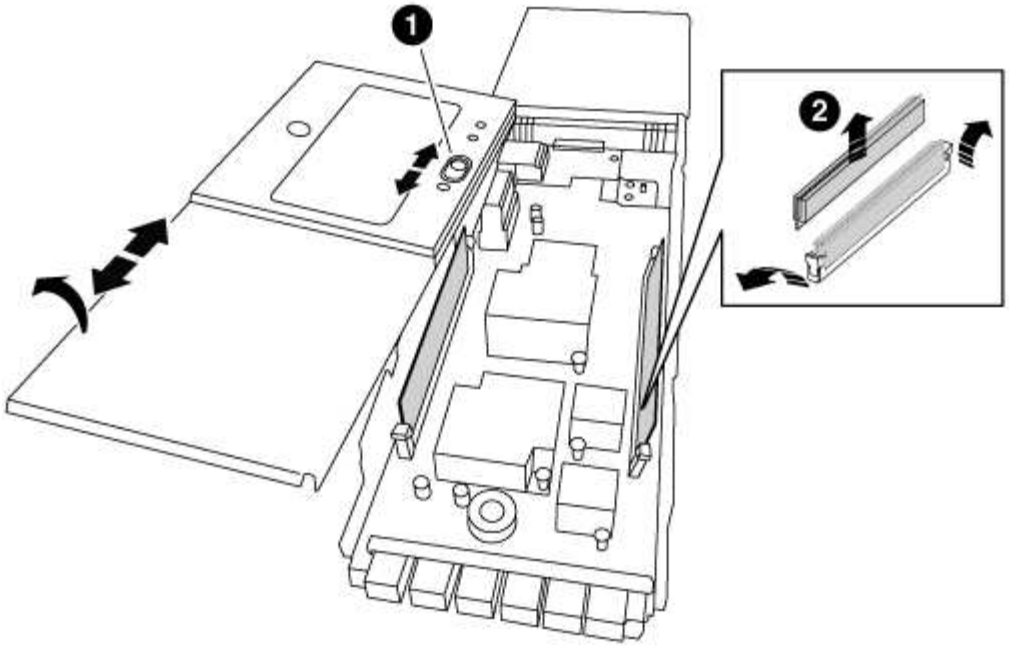

1	Cover locking button
2	DIMM and DIMM ejector tabs

4. Locate the DIMM to be replaced inside the NVRAM module, and then remove it by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
5. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
6. Close the cover on the module.
7. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

Step 4: Reboot the controller after FRU replacement

After you replace the FRU, you must reboot the controller module.

Step

1. To boot ONTAP from the LOADER prompt, enter `bye`.

Step 5: Reassign disks

Depending on whether you have an HA pair or two-node MetroCluster configuration, you must either verify the reassignment of disks to the new controller module or manually reassign the disks.

Select one of the following options for instructions on how to reassign disks to the new controller.

Option 1: Verify ID (HA pair)

Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* node and then verify that the change was implemented.

Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the replacement node is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the replacement node, boot the node, entering `y` if you are prompted to override the system ID due to a system ID mismatch.

   ```
   boot_ontap bye
   ```

   The node will reboot, if autoboot is set.

3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* node console and then, from the healthy node, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired node, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

   ```
   node1> `storage failover show`
                                       Takeover
   Node              Partner           Possible     State Description
   ------------      ------------      --------
   -------------------------------------
   node1             node2             false        System ID changed
   on partner (Old:
                                                     151759755, New:
   151759706), In takeover
   node2             node1             -            Waiting for
   giveback (HA mailboxes)
   ```

4. From the healthy node, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`

   c. Wait for the `savecore`command to complete before issuing the giveback.

You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

d. Return to the admin privilege level: `set -privilege admin`

5. Give back the node:

a. From the healthy node, give back the replaced node's storage: `storage failover giveback -ofnode replacement_node_name`

The *replacement* node takes back its storage and completes booting.

If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration Guide for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the `System ID changed on partner` message.

6. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* node should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home
ID Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

7. If the system is in a MetroCluster configuration, monitor the status of the node: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each node will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

8. If the node is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a node on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The *replacement* node is the current owner of the disks on the disaster site.

     Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

9. If your system is in a MetroCluster configuration, verify that each node is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

10. Verify that the expected volumes are present for each node: `vol show -node node-name`
11. If you disabled automatic takeover on reboot, enable it from the healthy node: `storage failover modify -node replacement-node-name -onreboot true`

Option 2: Reassign ID (MetroCluster config)

Reassign the system ID in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `metrocluster node show -fields node-systemid,dr-partner-systemid`

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

```
dr-group-id cluster         node                 node-systemid dr-
partner-systemid
 ----------- --------------------- --------------------
------------ -------------------
 1           Cluster_A             Node_A_1             536872914
118073209
 1           Cluster_B             Node_B_1             118073209
536872914
 2 entries were displayed.
```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

```
Local System ID: 118065481
    ...
    ...
```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                  POOL   SERIAL NUMBER  HOME
-------    -------------          -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Verify that the coredumps are saved: `system node run -node` *`local-node-name`* `partner savecore`

      If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node` *`local-node-name`* `partner savecore -s command`.</info>.

   c. Return to the admin privilege level: `set -privilege admin`
7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
8. Boot the *replacement* node: `boot_ontap`
9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`
10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

       You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

    b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`
    c. Return to the admin privilege level: `set -privilege admin`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - AFF A700

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- The number of power supplies in the system depends on the model.
- Power supplies are auto-ranging.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Press and hold the orange button on the power supply handle, and then pull the power supply out of the chassis.

> When removing a power supply, always use two hands to support its weight.

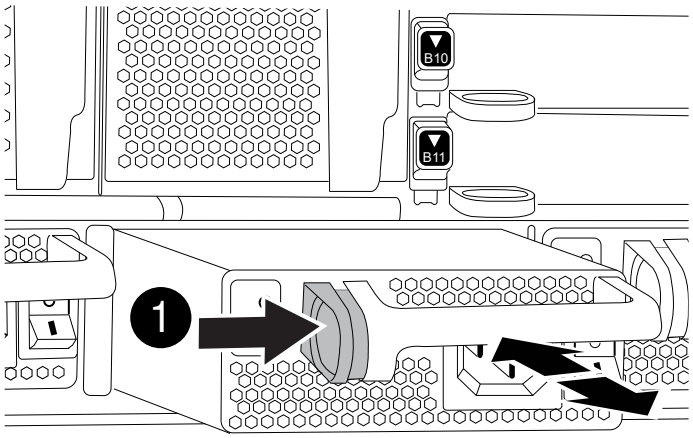

1	Locking button

5. Make sure that the on/off switch of the new power supply is in the Off position.

6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

   > Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply and the power source.

   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.

8. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

   The green power LED lights when the PSU is fully inserted into the chassis and the amber attention LED flashes initially, but turns off after a few moments.

9. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF 700

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the orange button on the cam handle downward until it unlocks.

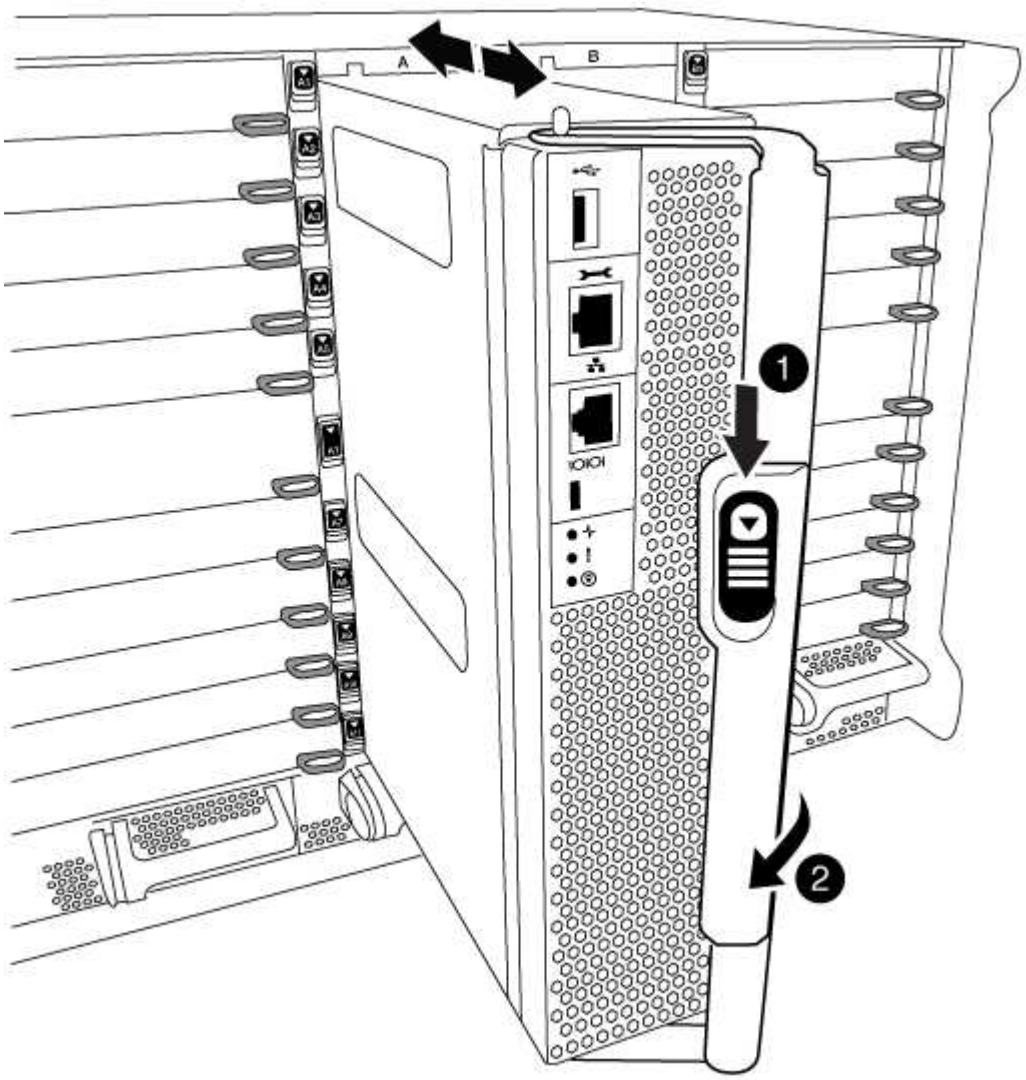

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

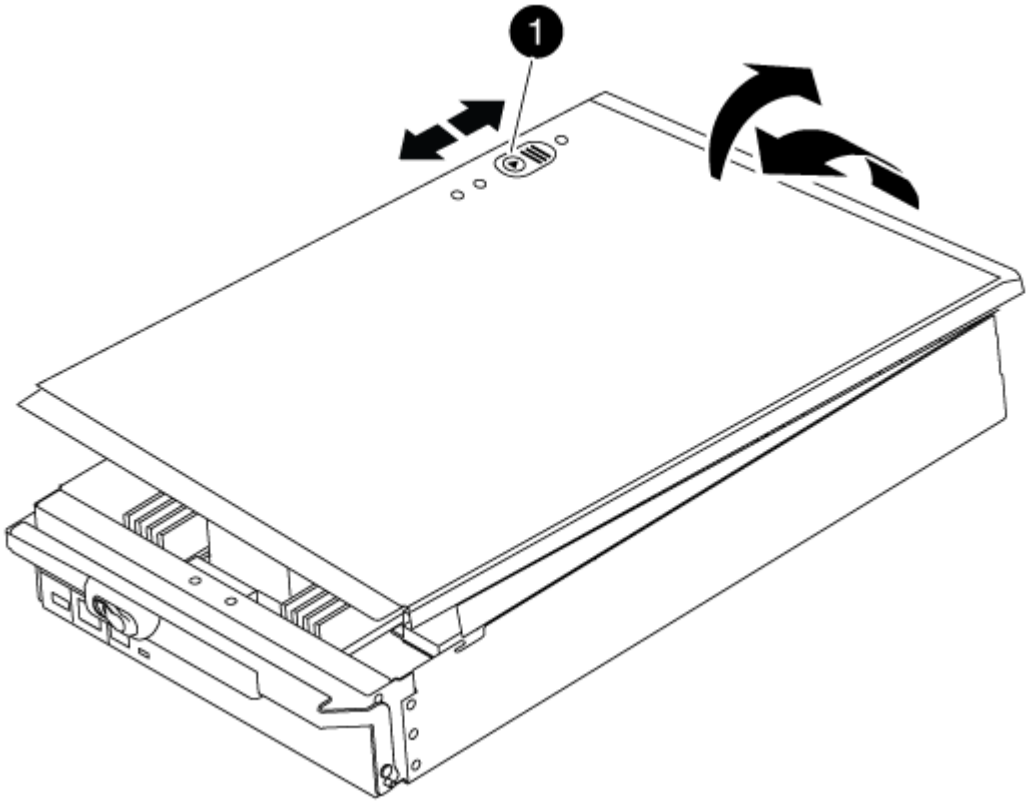

1	Controller module cover locking button

Step 3: Replace the RTC battery

To replace the RTC battery, you must locate the failed battery in the controller module, remove it from the holder, and then install the replacement battery in the holder.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

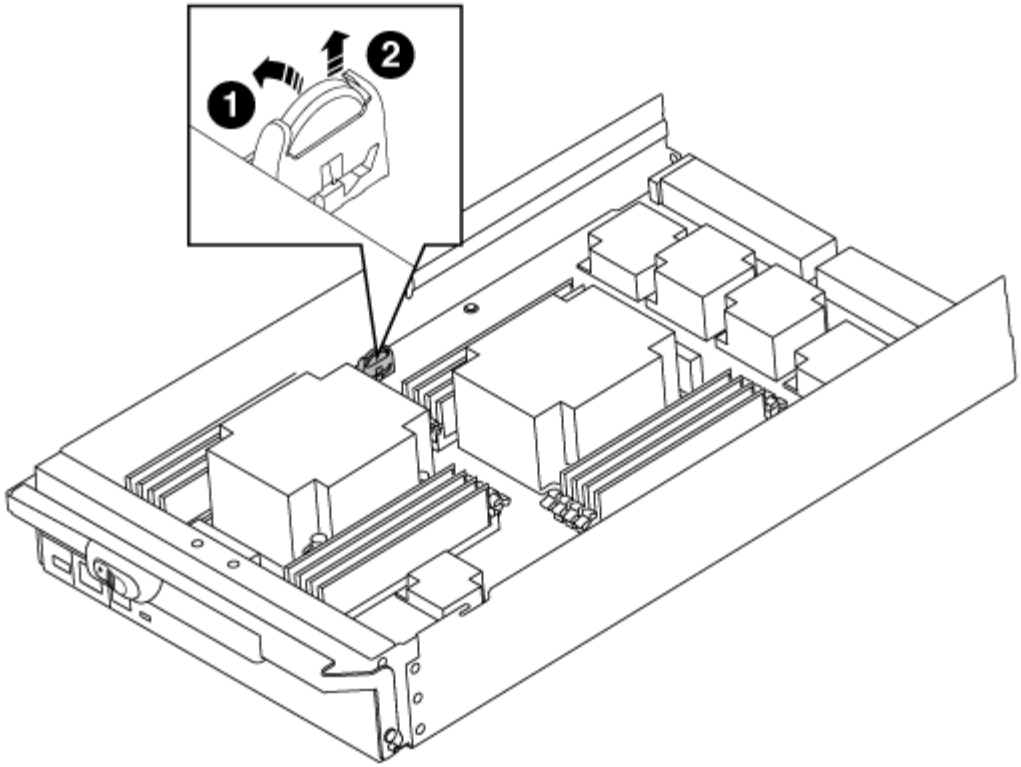

1	RTC battery
2	RTC battery housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
8. Reinstall the controller module cover.

Step 4: Reinstall the controller module and set time/date

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

Steps

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy node with the `show date` command.
   b. At the LOADER prompt on the target node, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target node.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the node reboot.
8. Return the node to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

X91148A module

Overview of adding an X91148A module - AFF A700

You can add an I/O module to your system by either replacing a NIC or storage adapter with a new one in a fully-populated system, or by adding a new NIC or storage adapter into an empty chassis slot in your system.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- To non-disruptively add an I/O module, you must takeover the target controller, remove the slot blanking cover in the target slot or remove an existing I/O module, add the new or replacement I/O module, and then giveback the target controller.
- Make sure that all other components are functioning properly.

Add an X91148A module in a system with open slots - A700

You can add an X91148A module into an empty module slot in your system as either a 100GbE NIC or a storage module for the NS224 storage shelves.

- Your system must be running ONTAP 9.8 and later.
- To non-disruptively add the X91148A module, you must takeover the target controller, remove the slot blanking cover in the target slot, add the module, and then giveback the target controller.
- There must be one or more open slots available on your system.
- If multiple slots are available, install the module according to the slot priority matrix for the X91148A module in the NetApp Hardware Universe.
- If you are adding the X91148A module as a storage module, you must install the module slots 3 and/or 7.
- If you are adding the X91148A module as a 100GbE NIC, you can use any open slot. However, by default, slots 3 and 7 are set as storage slots. If you wish to use those slots as network slots and will not add NS224 shelves, you must modify the slots for networking use with the `storage port modify -node` *`node name`* `-port` *`port name`* `-mode network` command. See the NetApp Hardware Universe for other slots that can be used by the X91148A module for networking.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Option 1: Add an X91148A module as a NIC module

To add an X91148A module as a NIC module in a system with open slots, you must follow the specific sequence of steps.

Steps

1. Shutdown controller A:
   a. Disable automatic giveback: `storage failover modify -node local -auto-giveback false`
   b. Take over the target node: `storage failover takeover -ofnode target_node_name`

      The console connection shows that the node drops to the LOADER prompt when the takeover is complete.
2. If you are not already grounded, properly ground yourself.
3. Remove the target slot blanking cover:
   a. Depress the lettered and numbered cam button.
   b. Rotate the cam latch down until it is in a horizontal position.
   c. Remove the blanking cover.
4. Install the X91148A module:
   a. Align the X91148A module with the edges of the slot.
   b. Slide the X91148A module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
5. Cable the module to the data switches.
6. Reboot controller A from the LOADER prompt: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.
7. Giveback the node from the partner node: `storage failover giveback -ofnode target_node_name`
8. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto-giveback true`
9. Repeat these steps for controller B.

Option 2: Add an X91148A module as a storage module

To add an X91148A module as a storage module in a system with open slots, you must follow the specific sequence of steps.

- This procedure presumes slots 3 and/or 7 are open.

Steps

1. Shut down controller A:
   a. Disable automatic giveback: `storage failover modify -node local -auto-giveback false`

b. Take over the target node: `storage failover takeover -ofnode target_node_name`

   The console connection shows that the node drops to the LOADER prompt when the takeover is complete.

2. If you are not already grounded, properly ground yourself.
3. Remove the target slot blanking cover:
   a. Depress the lettered and numbered cam button.
   b. Rotate the cam latch down until it is in a horizontal position.
   c. Remove the blanking cover.
4. Install the X91148A module into slot 3:
   a. Align the X91148A module with the edges of the slot.
   b. Slide the X91148A module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
   d. If you are installing a second X91148A module for storage, repeat this step for the module in slot 7.
5. Reboot controller A:
   - If the replacement module is not the same model as the old module, reboot the BMC :
     a. From the LOADER prompt, change to advanced privilege mode: `set -privilege advanced`
     b. Reboot the BMC: `sp reboot`
   - If the replacement module is the same as the old module, boot from the LOADER prompt: `bye`

     This reinitializes the PCIe cards and other components and reboots the node.

6. Giveback the node from the partner node: `storage failover giveback -ofnode target_node_name`
7. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto-giveback true`
8. Repeat these steps for controller B.
9. Install and cable your NS224 shelves, as described in Hot-add workflow.

Add an X91148A storage module in a system with no open slots - A700

You must remove one more or more existing NIC or storage modules in your system in order to install one or more X91148A storage modules into your fully-populated system.

- Your system must be running ONTAP 9.8 and later.
- To non-disruptively add the X91148A module, you must takeover the target controller, add the module, and then giveback the target controller.
- If you are adding the X91148A module as a storage adapter, you must install the module in slots 3 and/or 7.

- If you are adding the X91148A module as a 100GbE NIC, you can use any open slot. However, by default, slots 3 and 7 are set as storage slots. If you wish to use those slots as network slots and will not add NS224 shelves, you must modify the slots for networking use with the `storage port modify -node` *`node name`* `-port` *`port name`* `-mode network` command for each port. See the NetApp Hardware Universe for other slots that can be used by the X91148A module for networking.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Option 1: Add an X91148A module as a NIC module

You must remove one or more existing NIC or storage modules in your system in order to install one or more X91148A NIC modules into your fully-populated system.

Steps

1. If you are adding an X91148A module into a slot that contains a NIC module with the same number of ports as the X91148A module, the LIFs will automatically migrate when its controller module is shut down. If the NIC module being replaced has more ports than the X91148A module, you must permanently reassign the affected LIFs to a different home port. See Migrating a LIF for information about using System Manager to permanently move the LIFs
2. Shut down controller A:
   a. Disable automatic giveback: `storage failover modify -node local -auto-giveback false`
   b. Take over the target node: `storage failover takeover -ofnode target_node_name`

      The console connection shows that the node drops to the LOADER prompt when the takeover is complete.
3. If you are not already grounded, properly ground yourself.
4. Unplug any cabling on the target I/O module.
5. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.
   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.
   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

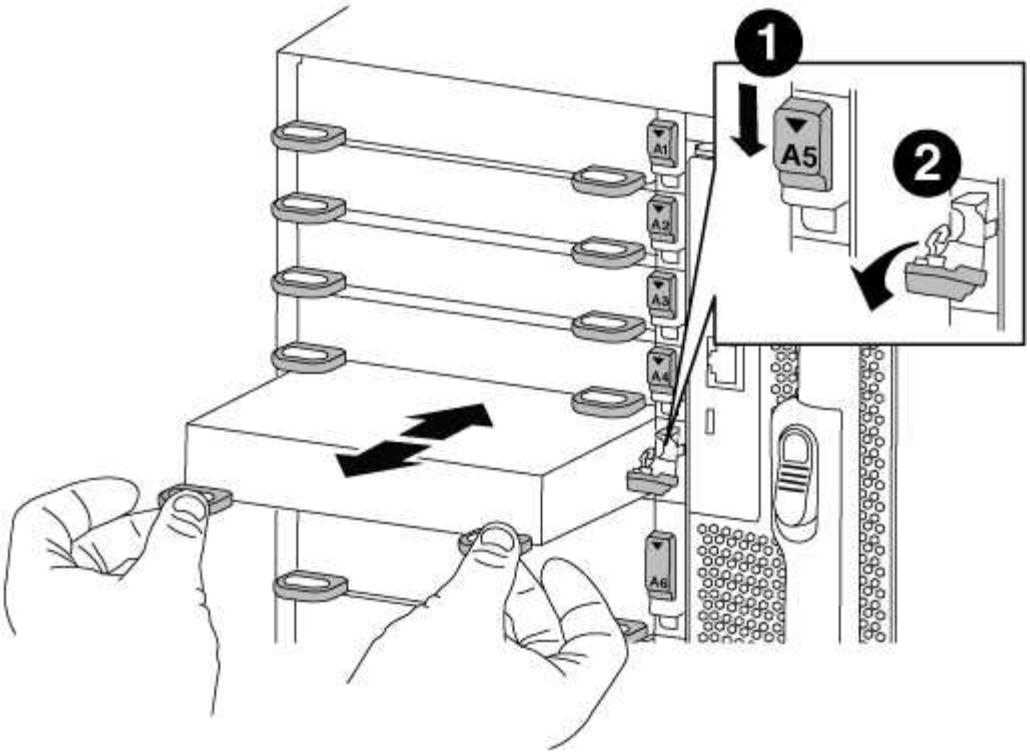

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

6. Install the X91148A module into the target slot:
   a. Align the X91148A module with the edges of the slot.
   b. Slide the X91148A module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
7. Repeat the remove and install steps to replace additional modules for controller A.
8. Cable the module or modules to the data switches.
9. Reboot controller A from the LOADER prompt: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.

10. Giveback the node from the partner node: `storage failover giveback -ofnode target_node_name`
11. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto -giveback true`
12. If you added the X91148A module as a NIC module in slots 3 or 7, for networking, use the `storage port modify -node` *node name* `-port` *port name* `-mode network` command for each port.

13. Repeat these steps for controller B.

Option 2: Adding an X91148A module as a storage module

You must remove one or more existing NIC or storage modules in your system in order to install one or more X91148A storage modules into your fully-populated system.

- This procedure presumes you re installing the X91148A module into slots 3 and/or 7.

Steps

1. If you are adding an X91148A module as a storage module in slots 3 and/or 7 into a slot that has an existing NIC module in it, use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.
2. Shut down controller A:
   a. Disable automatic giveback: `storage failover modify -node local -auto-giveback false`
   b. Take over the target node: `storage failover takeover -ofnode target_node_name`

      The console connection shows that the node drops to the LOADER prompt when the takeover is complete.
3. If you are not already grounded, properly ground yourself.
4. Unplug any cabling on the target I/O module.
5. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.
   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.
   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

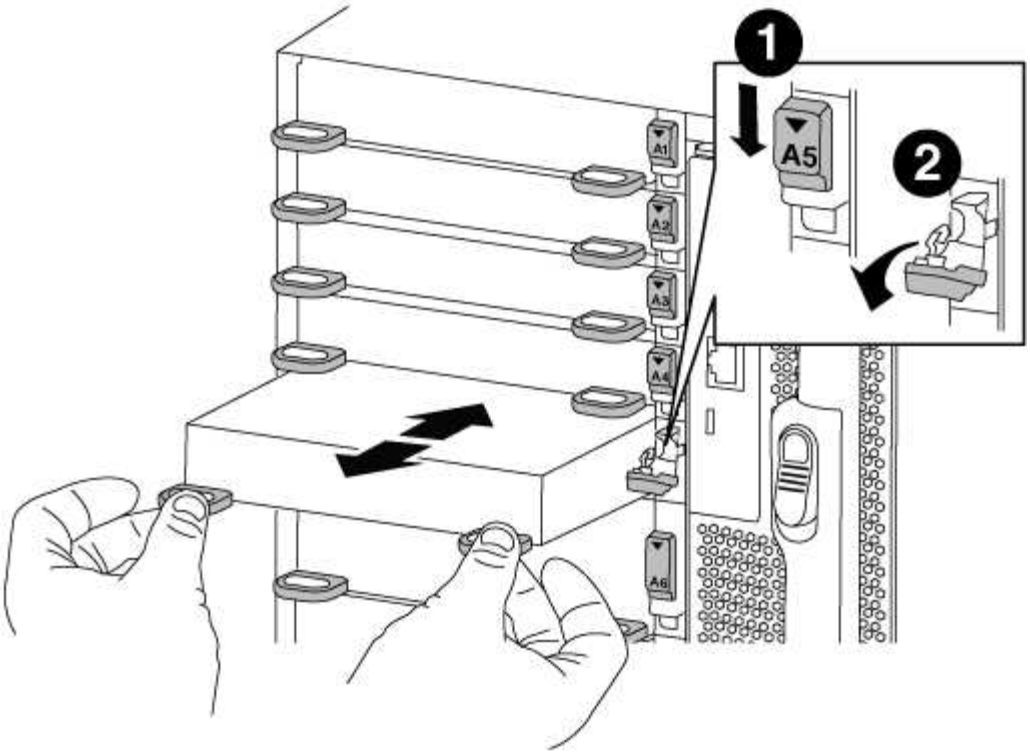

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

6. Install the X91148A module into slot 3:
   a. Align the X91148A module with the edges of the slot.
   b. Slide the X91148A module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
   d. If you are installing a second X91148A module for storage, repeat the remove and install steps for the module in slot 7.
7. Reboot controller A from the LOADER prompt: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.

8. Giveback the node from the partner node: `storage failover giveback -ofnode` *`target_node_name`*
9. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto -giveback true`
10. Repeat these steps for controller B.
11. Install and cable your NS224 shelves, as described in Hot-add workflow.

AFF A700s systems

Install and setup

Cluster configuration worksheet - AFF A700s

You can use the worksheet to gather and record your site-specific IP addresses and other information required when configuring an ONTAP cluster.

Cluster Configuration Worksheet

Start here: Choose your installation and setup experience

You can choose from different content formats to guide you through installing and setting up your new storage system.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

Installation and setup PDF poster - AFF A700s

You can use the PDF poster to install and set up your new system. The PDF poster provides step-by-step instructions with live links to additional content.

AFF A700s Installation and Setup Instructions

Installation and setup video - AFF A700s

The following video shows end-to-end software configuration for systems running ONTAP 9.2.

AFF A700s Setup Video

Maintain

Maintain AFF A700s hardware

For the AFF A700s storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

Fan

The fan cools the controller.

NVRAM battery

A battery is included with the controller and provides the backup power if the AC power fails.

NVRAM module

The NVRAM module (non-volatile random-access memory) preserves cached data if the power fails.

PCIe card

A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real-time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF A700s

The primary boot media stores the ONTAP boot image that the system uses when it boots. You can restore the primary boot media image by using the ONTAP image on the secondary boot media, or if necessary, by using a USB flash drive.

If your secondary boot media has failed or is missing the image.tgz file, you must restore the primary boot media using a USB flash drive. The drive must be formatted to FAT32 and must have the appropriate amount of storage to hold the image_xxx.tgz file.

- The replacement process restores the var file system from the secondary boot media or USB flash drive to the primary boot media.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
  - The *impaired* controller is the controller on which you are performing maintenance.
  - The *healthy* controller is the HA partner of the impaired controller.

If you need to replace the secondary boot media while the primary boot media is installed and healthy, contact NetApp Support and mention the How to replace the secondary boot device of an AFF A700s KB article.

Check encryption key support and status - AFF A700s

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller - AFF A700s

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller displays…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Replace the boot media - AFF A700s

You must remove the controller module from the chassis, open it, and then replace the failed boot media.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Unplug the controller module power supply from the source, and then unplug the cable from the power supply.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

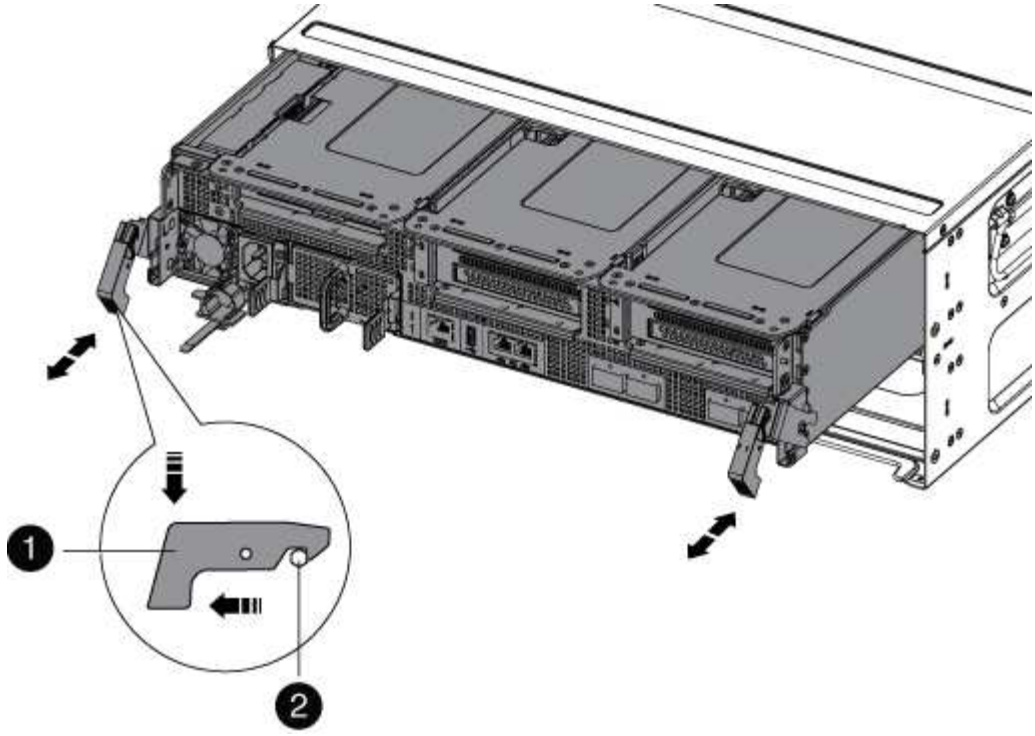

1	Locking latch
2	Locking pin

1. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

2. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

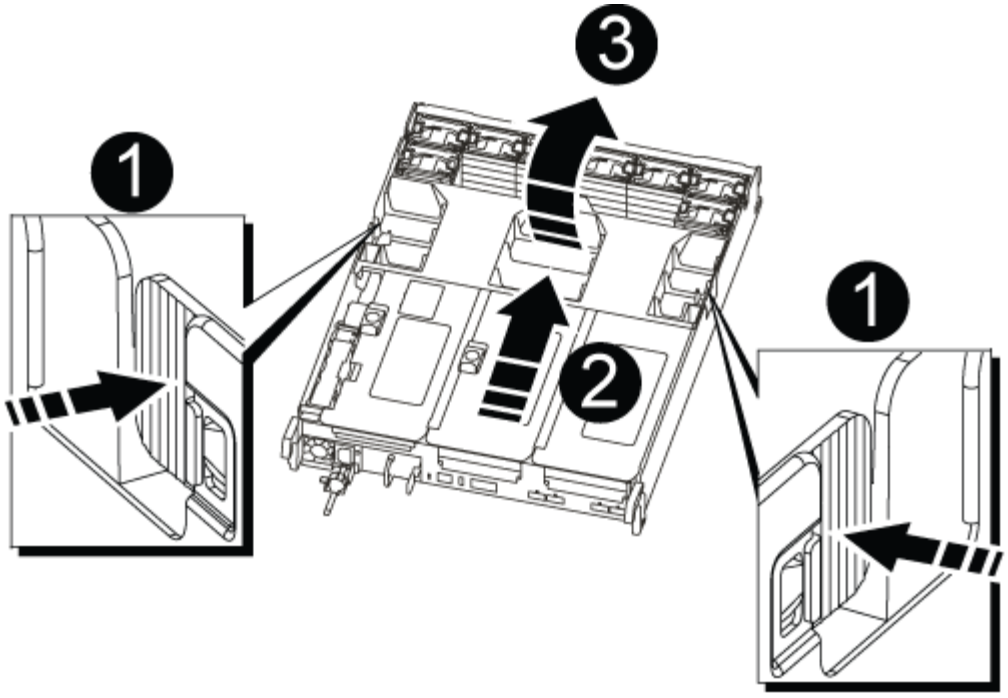

1	Air duct locking tabs
2	Risers
3	Air duct

Step 2: Replace the boot media - AFF A700s

You must locate the failed boot media in the controller module by removing the middle PCIe module on the controller module, locate the failed boot media, and then replace the boot media.

You need a Phillips head screwdriver to remove the screw that holds the boot media in place.

1. If you are not already grounded, properly ground yourself.
2. Locate the boot media:
   a. Open the air duct, if needed.
   b. If needed, remove Riser 2, the middle PCIe module, by unlocking the locking latch and then removing the riser from the controller module.

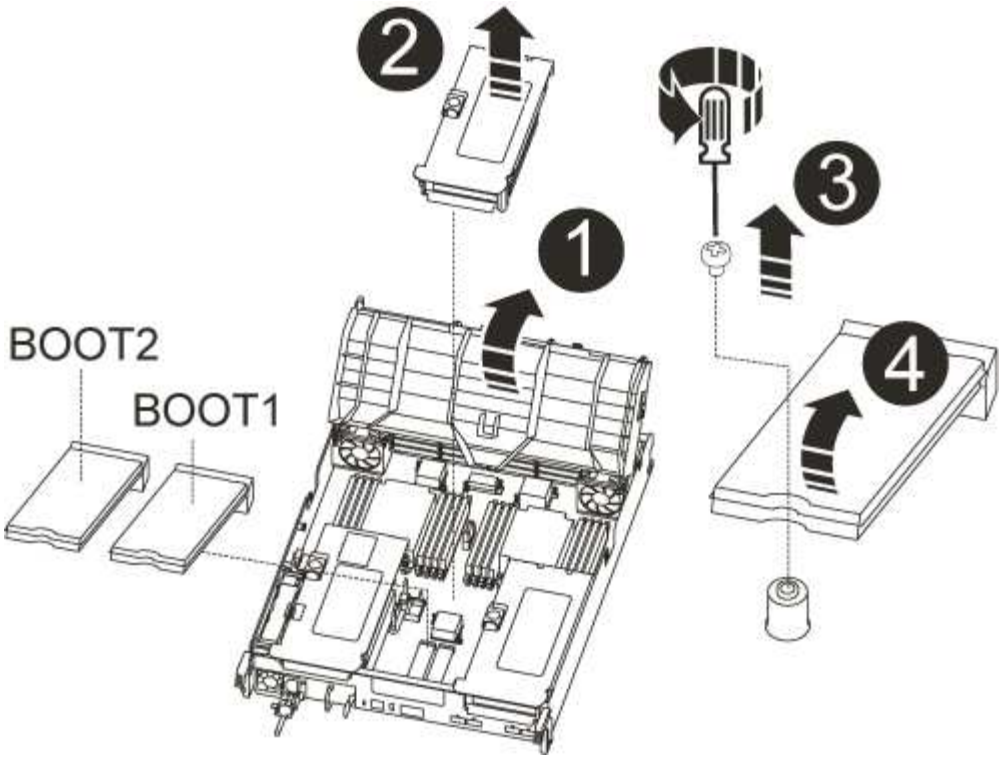

1	Air duct
2	Riser 2 (middle PCIe module)
3	Boot media screw
4	Boot media

3. Locate the failed boot media.
4. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.
   b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.
5. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
6. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.
7. Rotate the boot media down until it is flush with the motherboard.
8. Secure the boot media in place by using the screw.

Do not over-tighten the screw. Doing so might crack the boot media circuit board.

9. Reinstall the riser into the controller module.
10. Close the air duct:
    a. Rotate the air duct downward.
    b. Slide the air duct toward the risers until it clicks into place.

Transfer the boot image to the boot media - AFF A700s

You can install the system image to the replacement boot media using by using either the image on second boot media installed in the controller module, the primary method to restore the system image, or by transferring the boot image to the boot media using a USB flash drive when the secondary boot media restore failed or if the image.tgz file is not found on the secondary boot media.

Option 1: Transfer files to the boot media using backup recovery from the second boot media

You can install the system image to the replacement boot media using the image on second boot media installed in the controller module. This is the primary method for transferring the boot media files to the replacement boot media in systems with two boot media in the controller module.

The image on the secondary boot media must contain an `image.tgz` file and must not be reporting failures. If image.tgz file is missing or the boot media reports failures, you cannot use this procedure. You must transfer the boot image to the replacement boot media using the USB flash drive replacement procedure.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, close the air duct:
    a. Swing the air duct all the way down to the controller module.
    b. Slide the air duct toward the risers until the locking tabs click into place.
    c. Inspect the air duct to make sure that it is properly seated and locked into place.

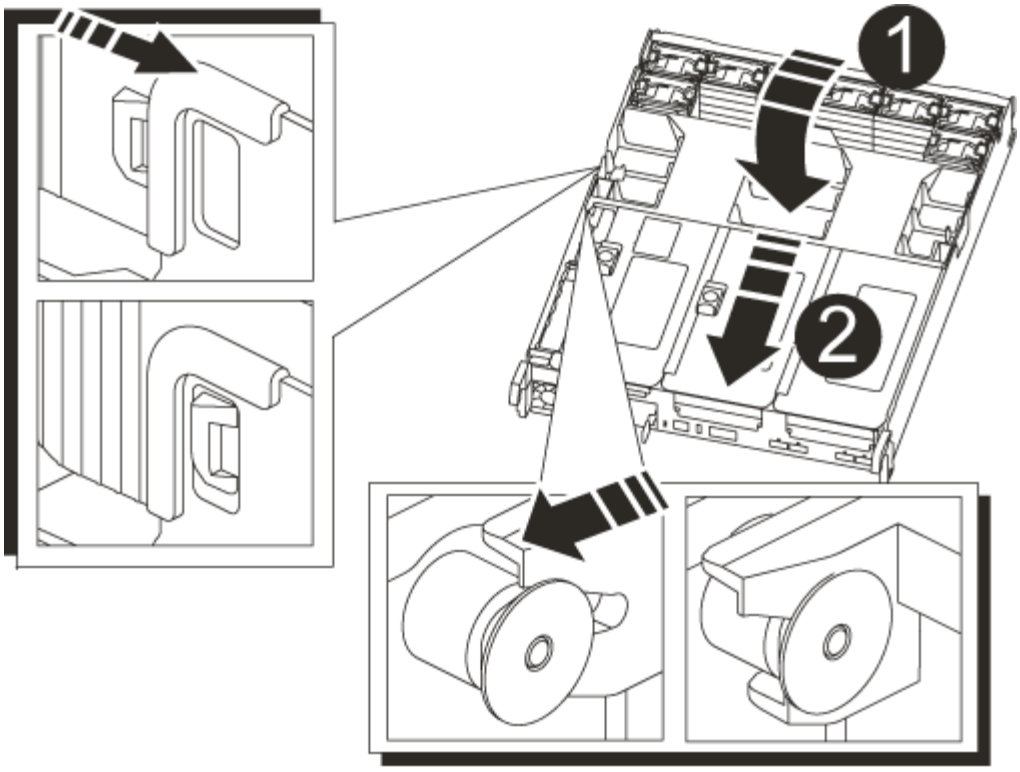

1	Air duct
2	Risers

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

5. Gently push the controller module all the way into the system until the controller module locking hooks begin to rise, firmly push on the locking hooks to finish seating the controller module, and then swing the locking hooks into the locked position over the pins on the controller module.
6. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

7. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

8. From the LOADER prompt, boot the recovery image from the secondary boot media: `boot_recovery`

   The image is downloaded from the secondary boot media.

9. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
10. After the image is installed, start the restoration process:
    a. Record the IP address of the impaired controller that is displayed on the screen.
    b. Press `y` when prompted to restore the backup configuration.
    c. Press `y` when prompted to confirm that the backup procedure was successful.
11. From the partner controller in advanced privilege level, start the configuration synchronization using the IP address recorded in the previous step: `system node restore-backup -node local -target -address impaired_node_IP_address`
12. After the configuration synchronization is complete without errors, press `y` when prompted to confirm that the backup procedure was successful.
13. Press `y` when prompted whether to use the restored copy, and then press `y` when prompted to reboot the controller.
14. Exit advanced privilege level on the healthy controller.

Option 2: Transfer the boot image to the boot media using a USB flash drive

This procedure should only be used if the secondary boot media restore failed or if the image.tgz file is not found on the secondary boot media.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, close the air duct:
    a. Swing the air duct all the way down to the controller module.
    b. Slide the air duct toward the risers until the locking tabs click into place.
    c. Inspect the air duct to make sure that it is properly seated and locked into place.

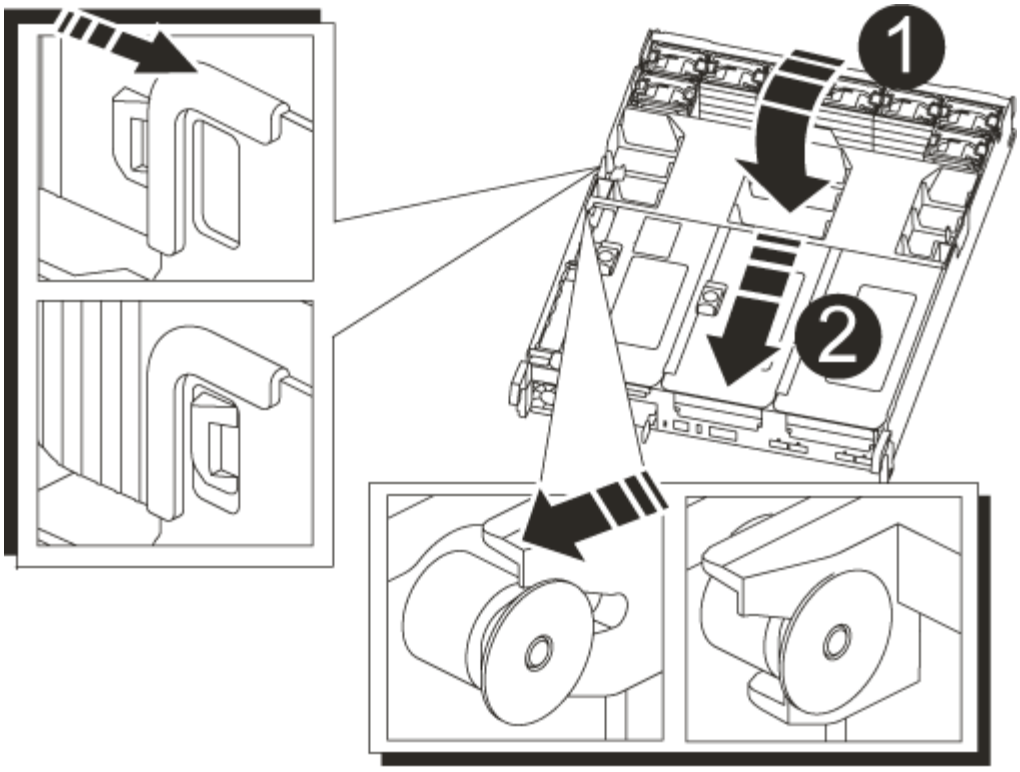

1	Air duct
2	Risers

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

5. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

6. Gently push the controller module all the way into the system until the controller module locking hooks begin to rise, firmly push on the locking hooks to finish seating the controller module, and then swing the locking hooks into the locked position over the pins on the controller module.
7. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

8. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

9. Although the environment variables and bootargs are retained, you should check that all required boot environment variables and bootargs are properly set for your system type and configuration using the `printenv bootarg name` command and correct any errors using the `setenv variable-name <value>` command.
   a. Check the boot environment variables:
      - `bootarg.init.boot_clustered`
      - `partner-sysid`
      - `bootarg.init.flash_optimized` for AFF C190/AFF A220 (All Flash FAS)
      - `bootarg.init.san_optimized` for AFF A220 and All-Flash SAN Array
      - `bootarg.init.switchless_cluster.enable`
   b. If External Key Manager is enabled, check the bootarg values, listed in the `kenv` ASUP output:
      - `bootarg.storageencryption.support <value>`
      - `bootarg.keymanager.support <value>`
      - `kmip.init.interface <value>`
      - `kmip.init.ipaddr <value>`
      - `kmip.init.netmask <value>`
      - `kmip.init.gateway <value>`
   c. If Onboard Key Manager is enabled, check the bootarg values, listed in the `kenv` ASUP output:
      - `bootarg.storageencryption.support <value>`
      - `bootarg.keymanager.support <value>`
      - `bootarg.onboard_keymanager <value>`
   d. Save the environment variables you changed with the `savenv` command
   e. Confirm your changes using the `printenv *variable-name*` command.
10. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

    The image is downloaded from the USB flash drive.
11. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
12. After the image is installed, start the restoration process:
    a. Record the IP address of the impaired controller that is displayed on the screen.
    b. Press `y` when prompted to restore the backup configuration.
    c. Press `y` when prompted to confirm that the backup procedure was successful.
13. Press `y` when prompted whether to use the restored copy, and then press `y` when prompted to reboot the controller.
14. From the partner controller in advanced privilege level, start the configuration synchronization using the IP address recorded in the previous step: `system node restore-backup -node local -target -address *impaired_node_IP_address*`

15. After the configuration synchronization is complete without errors, press `y` when prompted to confirm that the backup procedure was successful.
16. Press `y` when prompted whether to use the restored copy, and then press `y` when prompted to reboot the controller.
17. Verify that the environmental variables are set as expected.
    a. Take the controller to the LOADER prompt.

       From the ONTAP prompt, you can issue the command 'system node halt -skip-lif-migration-before -shutdown true -ignore-quorum-warnings true -inhibit-takeover true'.

    b. Check the environment variable settings with the `printenv` command.
    c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
    d. Save your changes using the `savenv` command.
    e. Reboot the controller.
18. With the rebooted impaired controller displaying the `Waiting for giveback…` message, perform a giveback from the healthy controller:

If your system is in…	Then…
An HA pair	After the impaired controller is displaying the `Waiting for giveback…` message, perform a giveback from the healthy controller: a. From the healthy controller: `storage failover giveback -ofnode partner_node_name` The impaired controller takes back its storage, finishes booting, and then reboots and is again taken over by the healthy controller. If the giveback is vetoed, you can consider overriding the vetoes. HA pair management b. Monitor the progress of the giveback operation by using the `storage failover show-giveback` command. c. After the giveback operation is complete, confirm that the HA pair is healthy and that takeover is possible by using the `storage failover show` command. d. Restore automatic giveback if you disabled it using the `storage failover modify` command.

19. Exit advanced privilege level on the healthy controller.

Boot the recovery image - AFF A700s

You must boot the ONTAP image from the USB drive, restore the file system, and verify

the environmental variables.

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:
   - If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed
   - If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.
6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.
8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore encryption - AFF A700s

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*******************************************************************
**************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*******************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF A700s

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF A700s

To replace the chassis, you must move the controller modules and SSD drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in its place.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.

- This procedure is written with the assumption that you are moving the SSDs and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A700s

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   `Warning: Are you sure you want to halt node <node_name>? {y|n}:`

10. Wait for each controller to halt and display the LOADER prompt.

Replace hardware - AFF A700s

Move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supply from the source, and then unplug the cable from the power supply.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

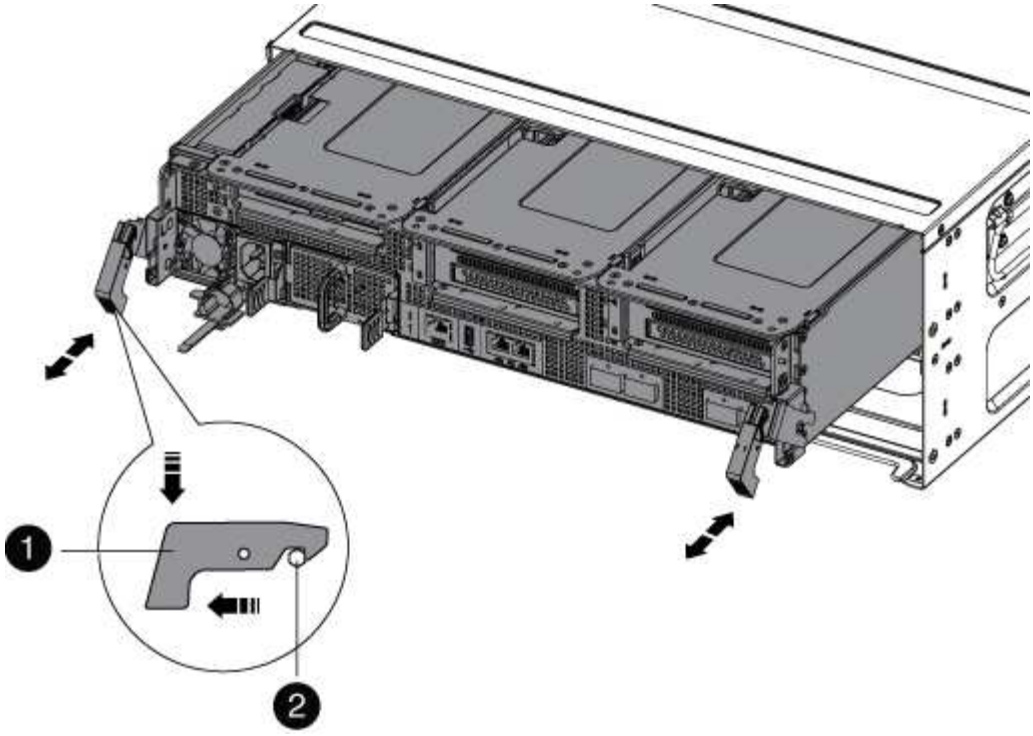

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.

3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.

5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It click when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controllers

After you install the controller module into the new chassis, boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. If you are not already grounded, properly ground yourself.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the console to the controller module, and then reconnect the management port.
4. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

d. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

e. Interrupt the boot process by pressing `Ctrl-C` when you see `Press Ctrl-C for Boot Menu`.

f. Select the option to boot to Maintenance mode from the displayed menu.

5. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - AFF A700s

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - AFF A700s

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - AFF A700s

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the controller module hardware - AFF A700s

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Unplug the controller module power supply from the source, and then unplug the cable from the power supply.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

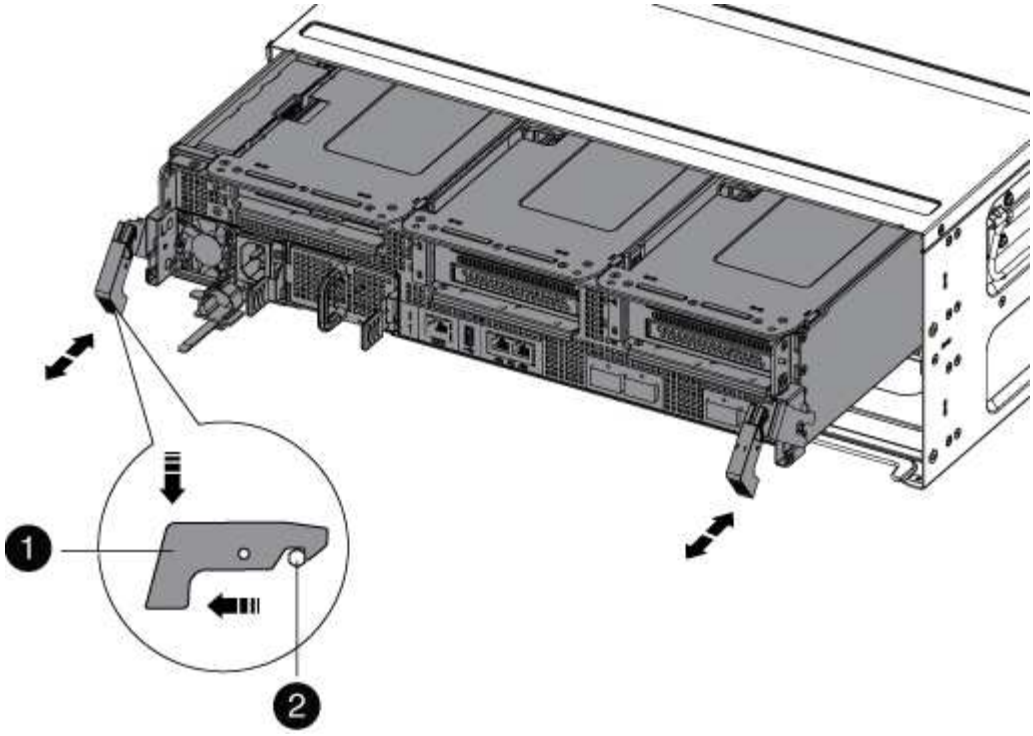

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

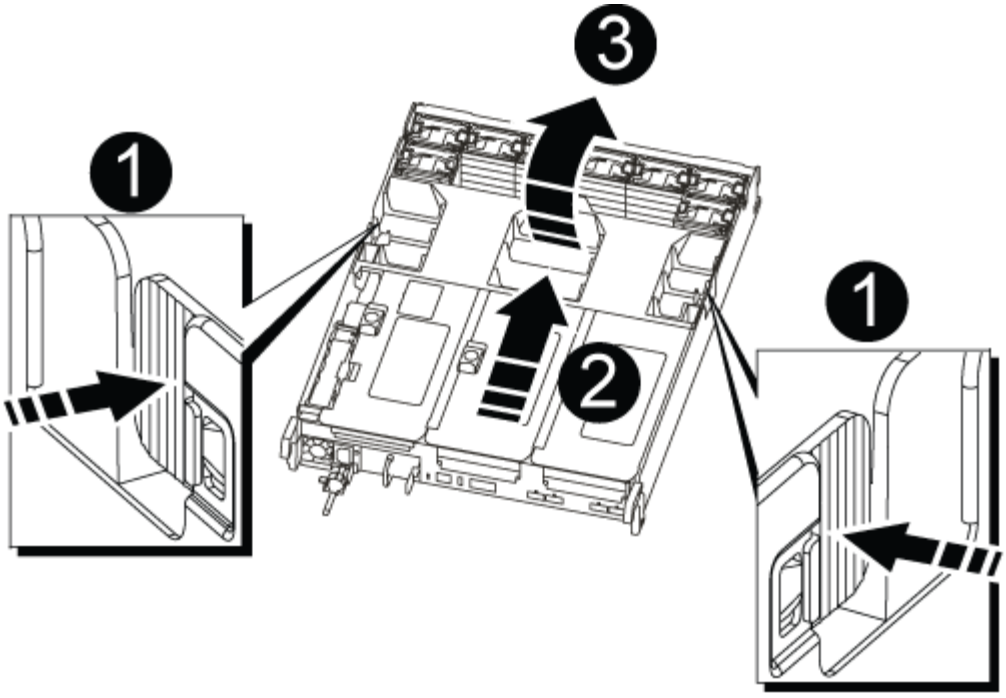

1	Air duct locking tabs
2	Risers
3	Air duct

Step 2: Move the NVRAM card

As part of the controller replacement process, you must remove the NVRAM card from Riser 1 in the impaired controller module and install the card into Riser 1 of the replacement controller module. You should only reinstall Riser 1 into the replacement controller module after you have moved the DIMMs from the impaired controller module to the replacement controller module.

1. Remove the NVRAM riser, Riser 1, from the controller module:
   a. Rotate the riser locking latch on the left side of the riser up and toward the fans.

      The NVRAM riser raises up slightly from the controller module.

   b. Lift the NVRAM riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser straight up out of the controller module, and then place it on a stable, flat surface so that you can access the NVRAM card.

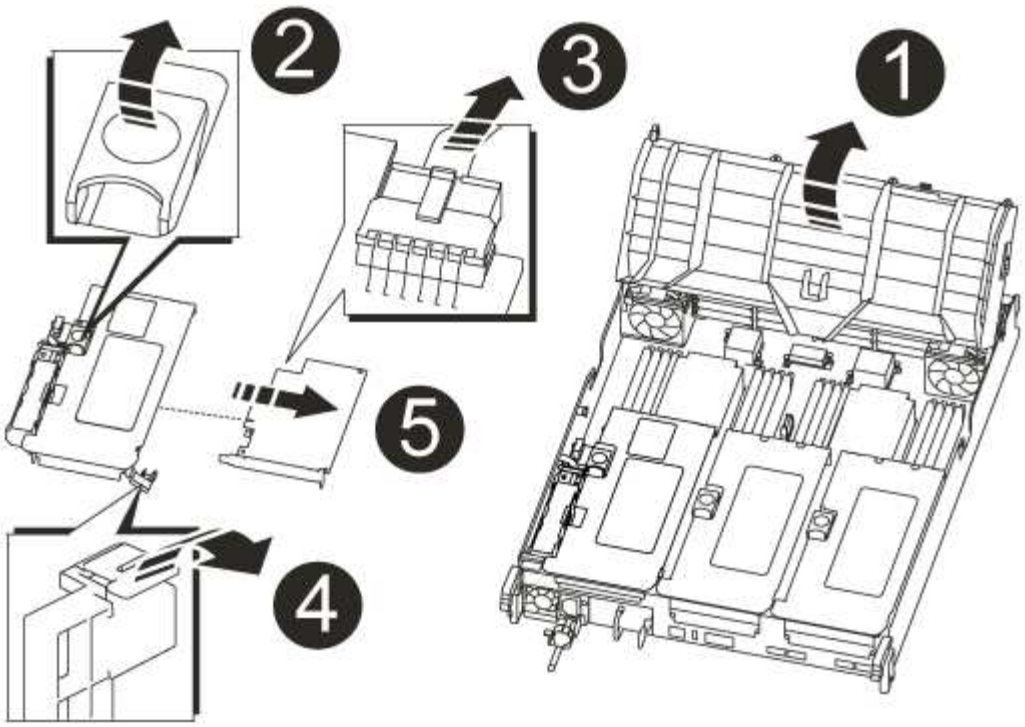

1	Air duct
2	Riser 1 locking latch
3	NVRAM battery cable plug connecting to the NVRAM card
4	Card locking bracket
5	NVRAM card

2. Remove the NVRAM card from the riser module:
   a. Turn the riser module so that you can access the NVRAM card.
   b. Unplug the NVRAM battery cable that is attached to the NVRAM card.
   c. Press the locking bracket on the side of the NVRAM riser, and then rotate it to the open position.
   d. Remove the NVRAM card from the riser module.
3. Remove the NVRAM riser from the replacement controller module.
4. Install the NVRAM card into the NVRAM riser:
   a. Align the card with the card guide on the riser module and the card socket in the riser.
   b. Slide the card squarely into the card socket.

   Make sure that the card is completely and squarely seated into the riser socket.

c. Connect the battery cable to the socket on the NVRAM card.

d. Swing the locking latch into the locked position and make sure that it locks in place.

Step 3: Move PCIe cards

As part of the controller replacement process, you must remove both PCIe riser modules, Riser 2 (the middle riser) and Riser 3 (riser on the far right) from the impaired controller module, remove the PCIe cards from the riser modules, and install the cards in the same riser modules in the replacement controller module. You will install the riser modules into the replacement controller module once the DIMMs have been moved to the replacement controller module.

Do not install the risers from the impaired controller module into the replacement controller module.

1. Remove the PCIe riser from the controller module:

   a. Remove any SFP modules that might be in the PCIe cards.

   b. Rotate the module locking latch on the left side of the riser up and toward the fan modules.

      The PCIe riser raises up slightly from the controller module.

   c. Lift the PCIe riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

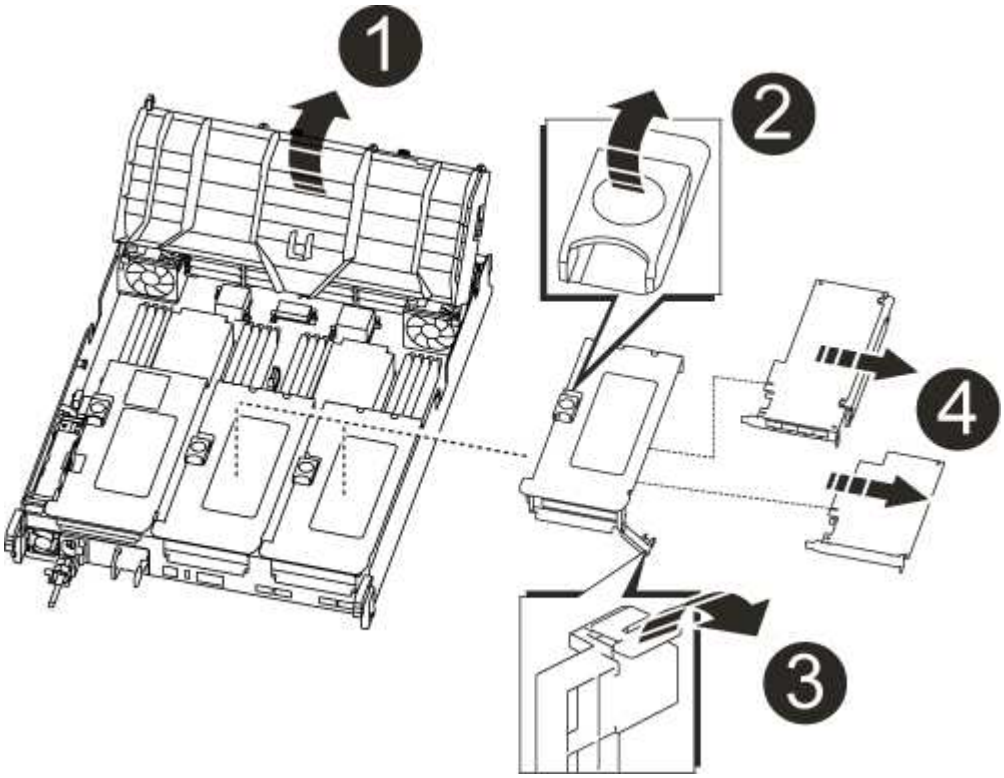

1	Air duct

2	Riser locking latch
3	Card locking bracket
4	Riser 2 (middle riser) and PCI cards in riser slots 2 and 3.

2. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. Remove the PCIe card from the riser.
3. Remove the corresponding riser from the replacement controller module.
4. Install the PCIe card into the riser from the replacement controller and then reinstall the riser back into the replacement controller:
   a. Align the card with the card guide on the riser and the card socket in the riser, and then slide it squarely into the socket in the riser.

      Make sure that the card is completely and squarely seated into the riser socket.

   b. Reinstall the riser into the replacement controller module.
   c. Swing the locking latch into place until it clicks into the locked position.
5. Repeat the preceding steps for Riser 3 and PCIe cards in slots 4 and 5 in the impaired controller module.

Step 4: Move the boot media

There are two boot media devices in the AFF A700s, a primary and a secondary or backup boot media. You must move them from the impaired controller to the *replacement* controller and install them into their respective slots in the *replacement* controller.

The boot media are located under Riser 2, the middle PCIe riser module. This PCIe module must be removed to gain access to the boot media.

1. Locate the boot media:
   a. Open the air duct, if needed.
   b. If needed, remove Riser 2, the middle PCIe module, by unlocking the locking latch and then removing the riser from the controller module.

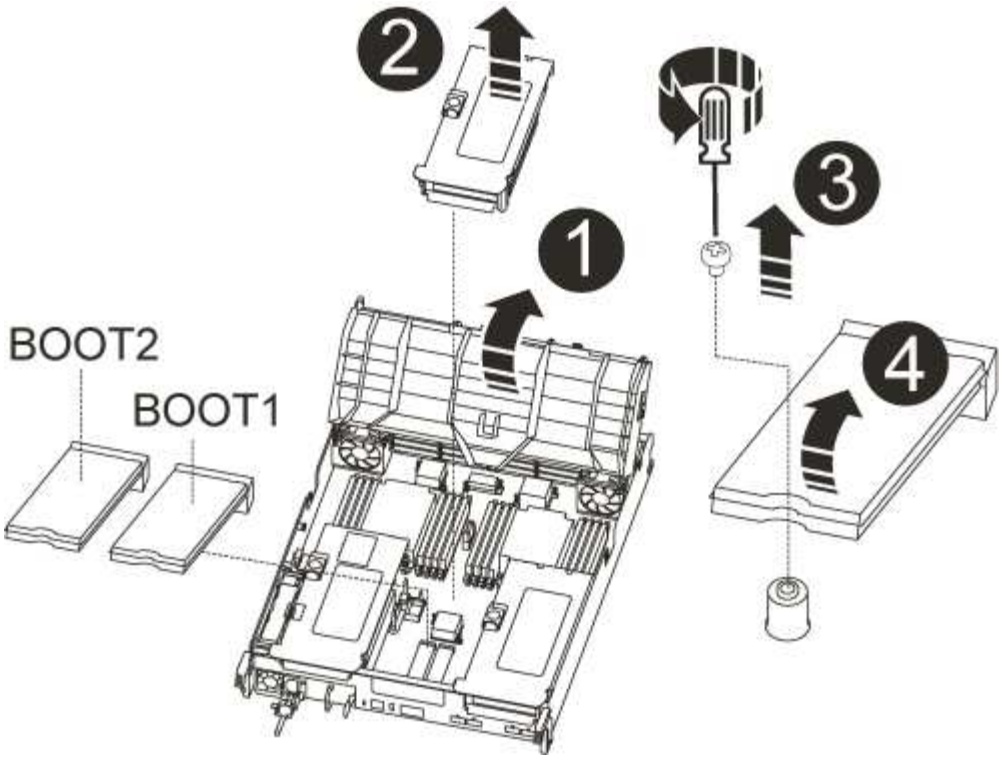

1	Air duct
2	Riser 2 (middle PCIe module)
3	Boot media screw
4	Boot media

2. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.
   b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.
3. Move the boot media to the new controller module and install it:

> Install the boot media into the same socket in the replacement controller module as it was installed in the impaired controller module; primary boot media socket (slot 1) to primary boot media socket, and secondary boot media socket (slot 2) to secondary boot media socket.

   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the motherboard.
   c. Secure the boot media to the motherboard using the boot media screw.

Do not over-tighten the screw or you might damage the boot media.

Step 5: Move the fans

You must move the fans from the impaired controller module to the replacement module when replacing a failed controller module.

1. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

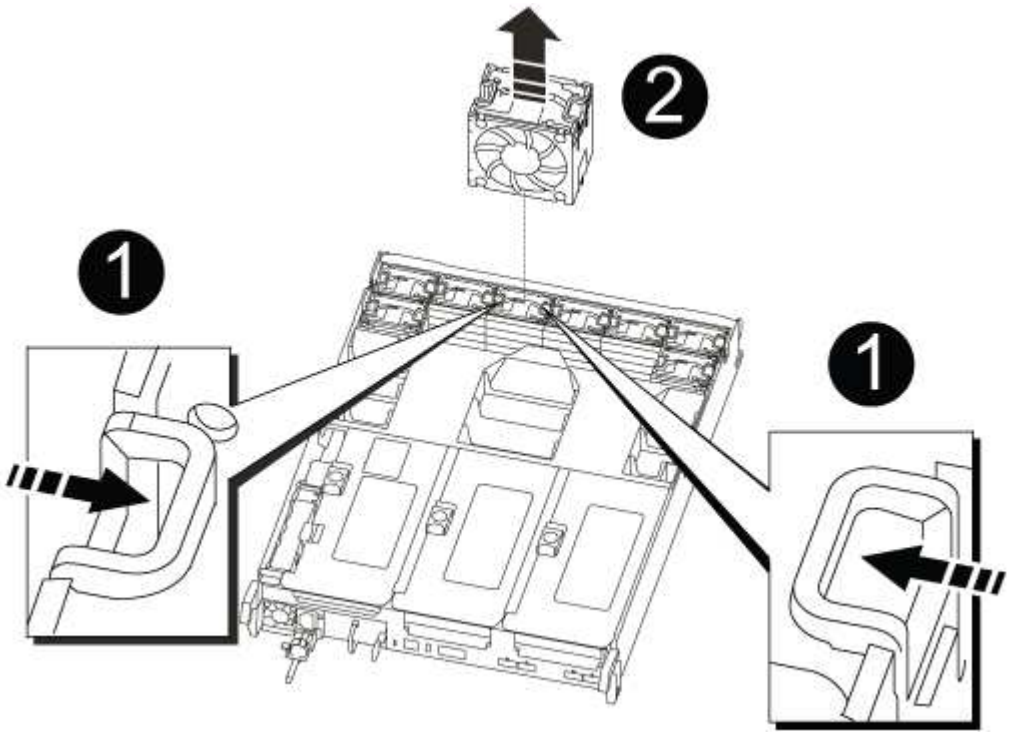

1	Fan locking tabs
2	Fan module

2. Move the fan module to the replacement controller module, and then install the fan module by aligning its edges with the opening in the controller module, and then sliding the fan module into the controller module until the locking latches click into place.
3. Repeat these steps for the remaining fan modules.

Step 6: Move system DIMMs

To move the DIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

1. Locate the DIMMs on your controller module.

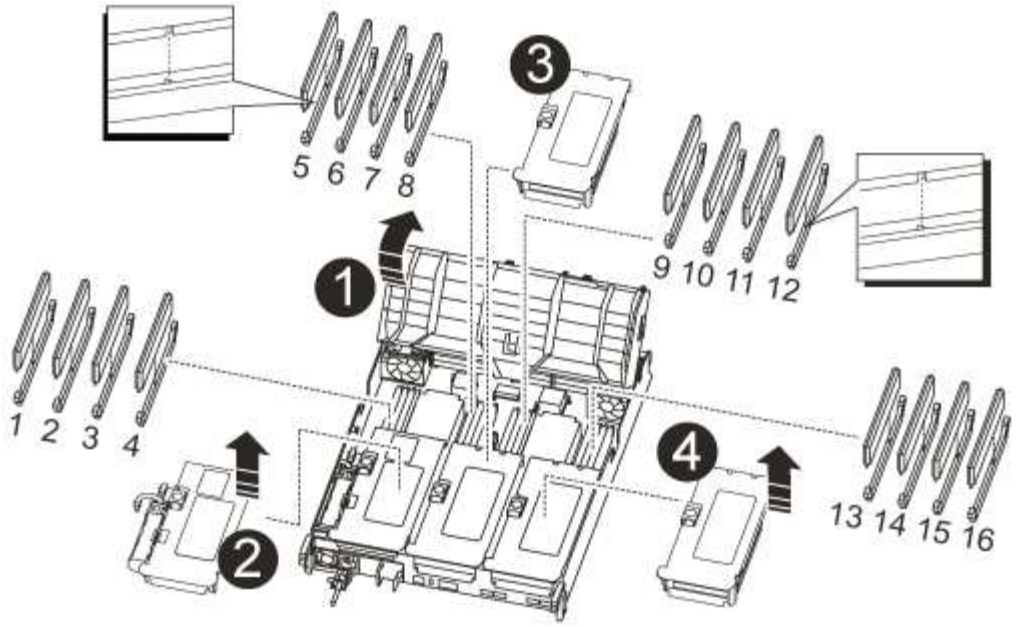

1	Air duct
2	Riser 1 and DIMM bank 1-4
3	Riser 2 and DIMM banks 5-8 and 9-12
4	Riser 3 and DIMM bank 13-16

2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

4. Locate the slot where you are installing the DIMM.
5. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Repeat these steps for the remaining DIMMs.

Step 7: Install the NVRAM module

To install the NVRAM module, you must follow the specific sequence of steps.

1. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

   d. Reinsert any SFP modules that were removed from the PCIe cards.

Step 8: Move the NVRAM battery

When replacing the controller module, you must move the NVRAM battery from the impaired controller module to the replacement controller module

1. Locate the NVRAM battery on the left side of the riser module, Riser 1.

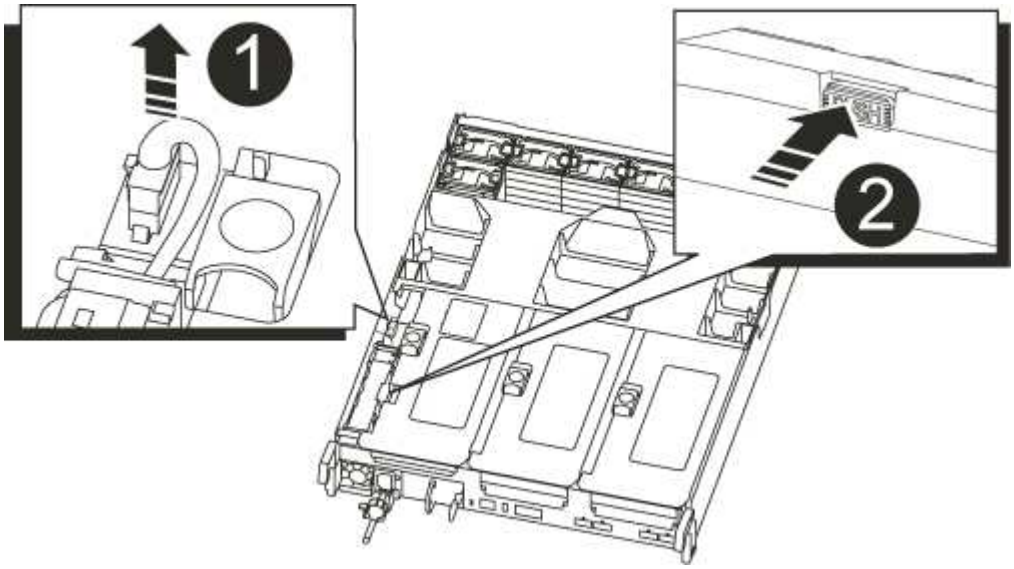

1	NVRAM battery plug
2	Blue NVRAM battery locking tab

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
4. Move the battery pack to the replacement controller module, and then install it in the NVRAM riser:
   a. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook

into the slots on the battery pack, and the battery pack latch engages and locks into place.

b. Press firmly down on the battery pack to make sure that it is locked into place.

c. Plug the battery plug into the riser socket and make sure that the plug locks into place.

Step 9: Install a PCIe riser

To install a PCIe riser, you must follow a specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

   d. Reinsert any SFP modules that were removed from the PCIe cards.
3. Repeat the preceding steps for Riser 3 and PCIe cards in slots 4 and 5 in the impaired controller module.

Step 10: Move the power supply

You must move the power supply and power supply blank from the impaired controller module to the replacement controller module when you replace a controller module.

1. If you are not already grounded, properly ground yourself.
2. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

> The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

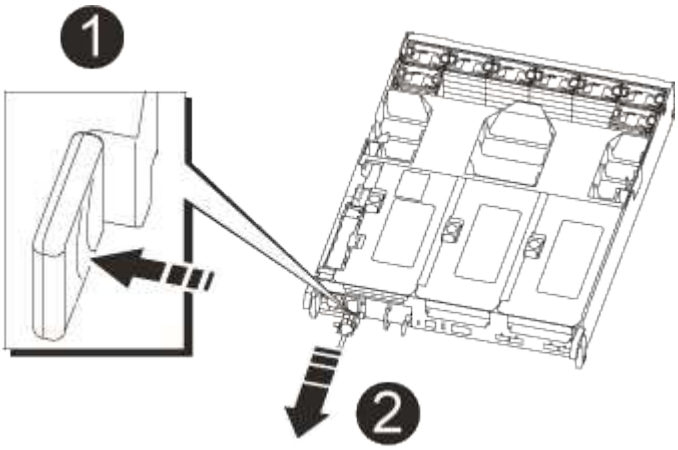

1	Blue power supply locking tab

2	Power supply

3. Move the power supply to the new controller module, and then install it.
4. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

5. Remove the PSU blanking panel from the impaired controller module, and then install it in the replacement controller module.

Step 11: Install the controller module

After all the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis and then boot it to Maintenance mode.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

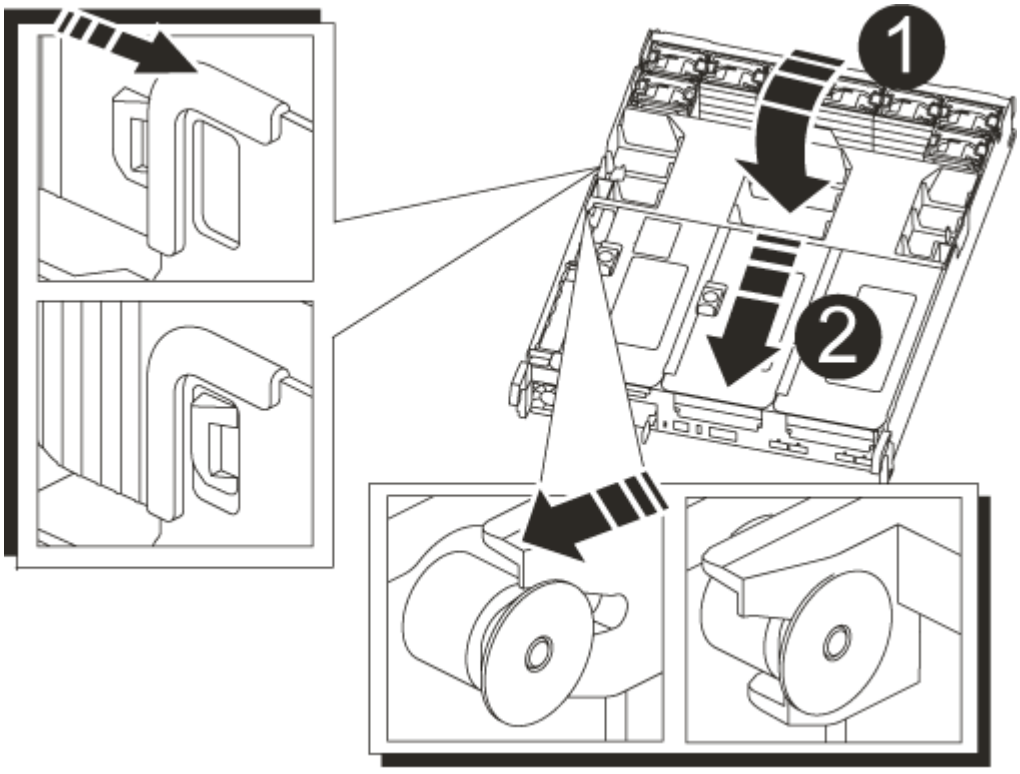

1	Locking tabs
2	Slide plunger

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. Interrupt the boot process by pressing `Ctrl-C`.
6. Plug the system cables and transceiver modules into the controller module and reinstall the cable management device.
7. Plug the power cables into the power supplies and reinstall the power cable retainers.
8. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs or onboard ports, convert these ports to 10 GbE connections by using the nicadmin convert command from Maintenance mode.

   Be sure to exit Maintenance mode after completing the conversion.

Restore and verify the system configuration - AFF A700s

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the chassis

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - ha
   - non-ha

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

To complete the replacement procedure and restore your system to full operation, you must recable the storage, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node                Partner         Possible     State Description
------------      ------------      --------
-------------------------------------
node1               node2             false        System ID changed on
partner (Old:
                                                    151759755, New:
151759706), In takeover
node2               node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed

on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Complete system restoration - AFF A700s

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement node in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF A700s

You must replace a DIMM in the controller when your storage system encounters errors

such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller.

Synchronize a node with the cluster

Steps

1. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
2. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows `Waiting for giveback…`, press Ctrl-C, and then respond `y`.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Unplug the controller module power supply from the source, and then unplug the cable from the power supply.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

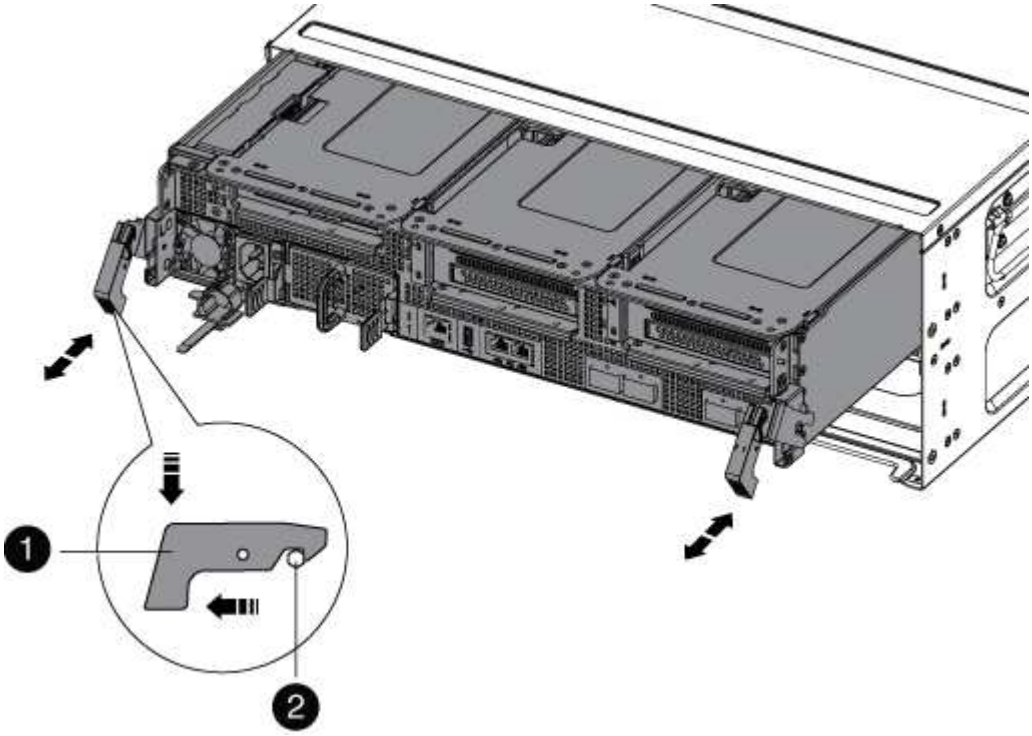

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

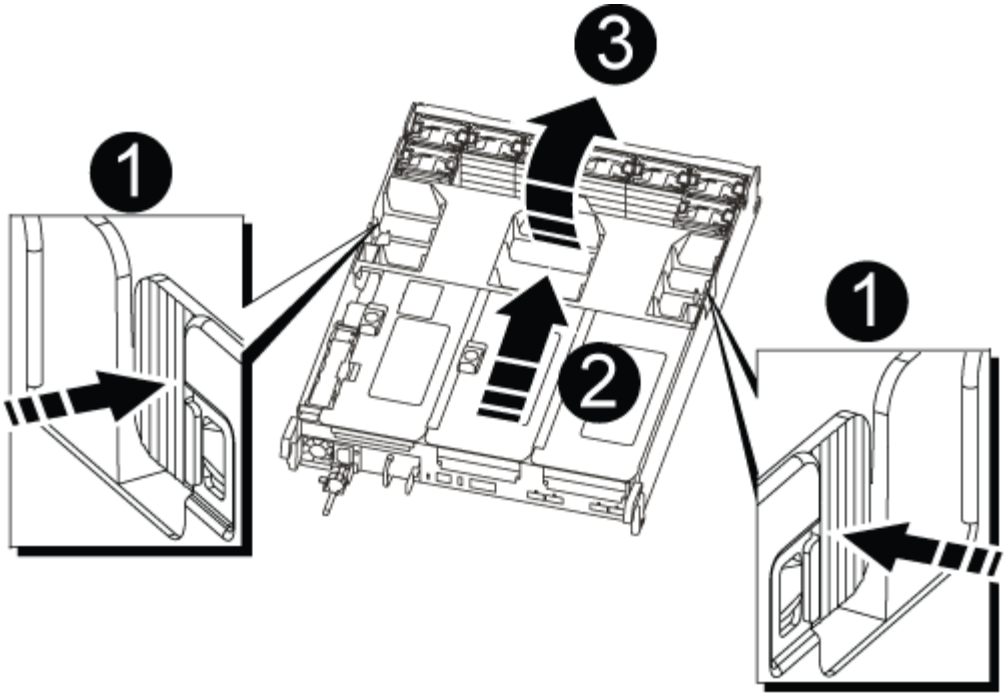

1	Air duct locking tabs
2	Risers
3	Air duct

Step 3: Replace a DIMM

To replace a DIMM, you must locate it in the controller module using the DIMM map on the inside of the controller module, and then replace it following the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the applicable riser.

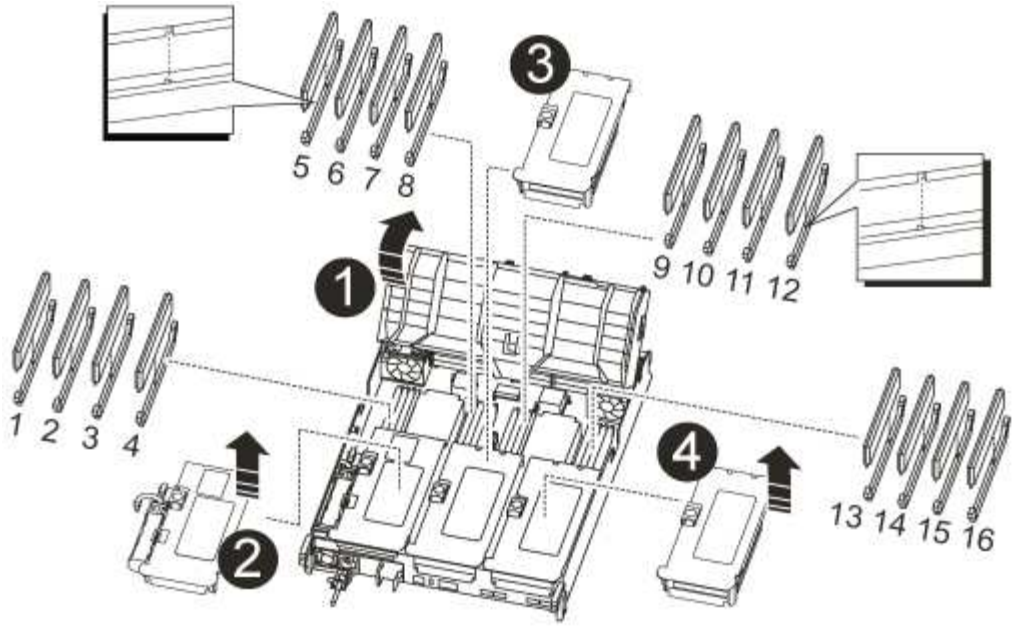

1	Air duct cover
2	Riser 1 and DIMM bank 1-4
3	Riser 2 and DIMM bank 5-8 and 9-12
4	Riser 3 and DIMM 13-16

- If you are removing or moving a DIMM in bank 1-4, unplug the NVRAM battery, unlock the locking latch on Riser 1, and then remove the riser.
- If you are removing or moving a DIMM in bank 5-8 or 9-12, unlock the locking latch on Riser 2, and then remove the riser.
- If you are removing or moving a DIMM in bank 13-16, unlock the locking latch on Riser 3, and then remove the riser.

3. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM

squarely into the slot.

The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Reinstall any risers that you removed from the controller module.

   If you removed the NVRAM riser, Riser 1, make sure that you plug the NVRAM battery into the controller module.

9. Close the air duct.

Step 4: Reinstall the controller module and boot the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. If you are not already grounded, properly ground yourself.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   d. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

Step 5:Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - AFF A700s

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  > Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan - AFF A800

To replace a fan, remove the failed fan module and replace it with a new fan module.

Step 1: Shut down the impaired controller - AFF A700s

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove the controller module - AFF A700s

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Unplug the controller module power supply from the source, and then unplug the cable from the power supply.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

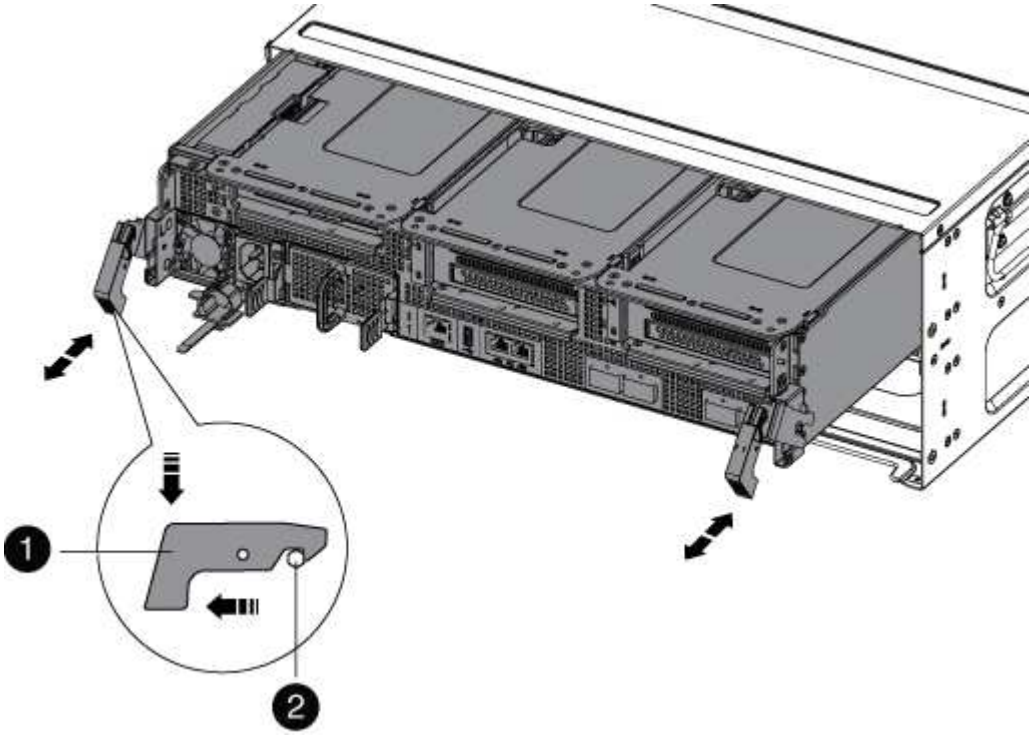

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface, and then open the air duct:

a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.

b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

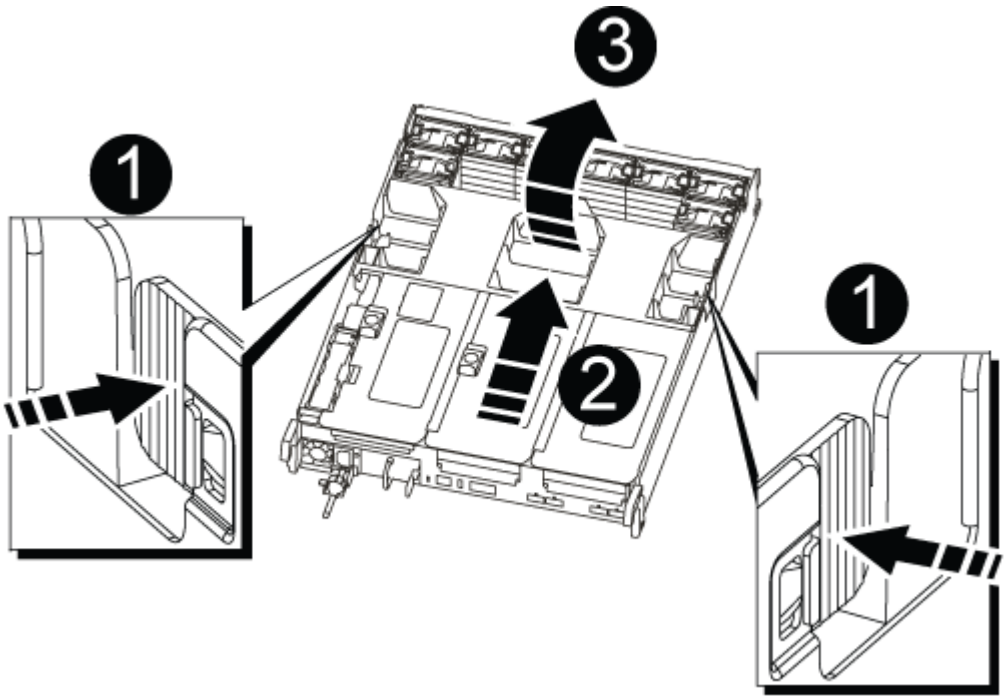

1	Air duct locking tabs
2	Risers
3	Air duct

Step 3: Replace the fan - AFF A700s

To replace a fan, remove the failed fan module and replace it with a new fan module.

1. If you are not already grounded, properly ground yourself.
2. Identify the fan module that you must replace by checking the console error messages.
3. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

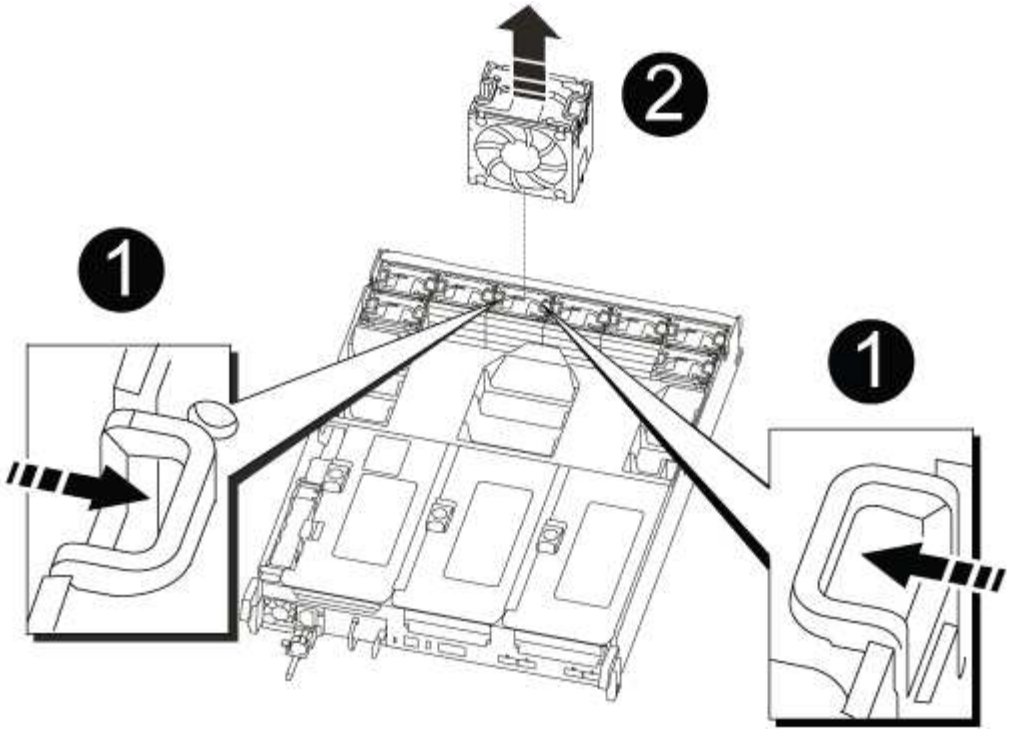

1	Fan locking tabs
2	Fan module

4. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module until the locking latches click into place.

Step 4: Reinstall the controller module - AFF A700s

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

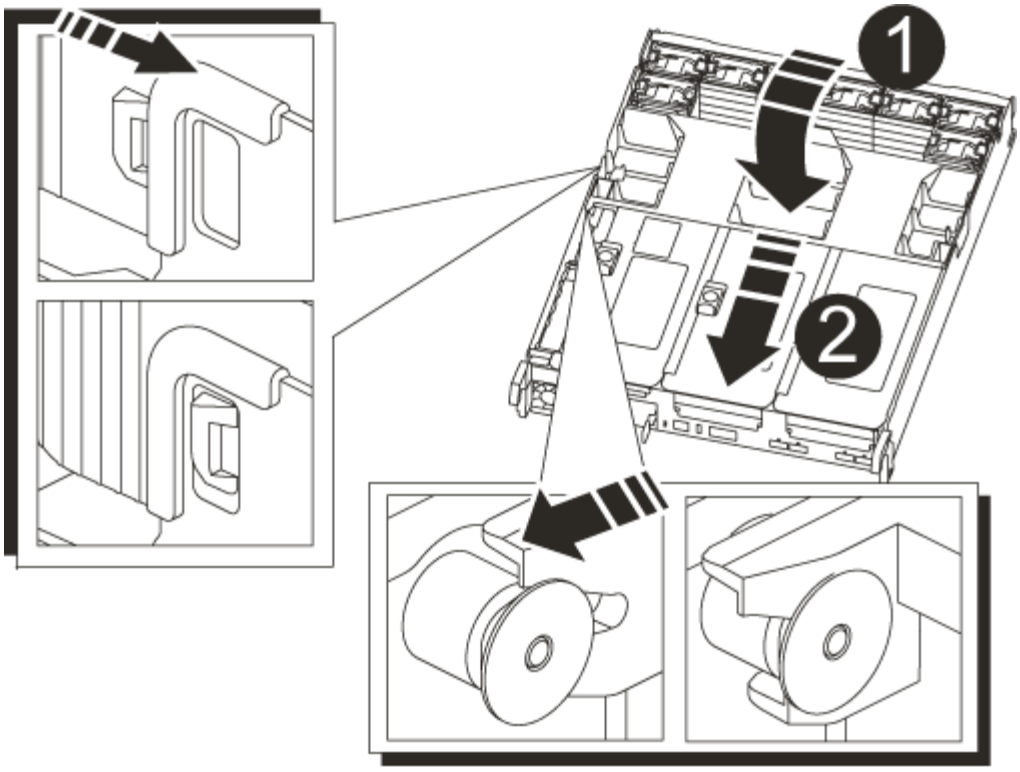

1	Locking tabs
2	Slide plunger

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   d. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect

the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

6. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs or onboard ports, convert these ports to 10 GbE connections by using the nicadmin convert command from Maintenance mode.

   Be sure to exit Maintenance mode after completing the conversion.

7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 5: Return the failed part to NetApp - AFF A700s

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVRAM battery - AFF A700s

To replace an NVRAM battery in the system, you must remove the controller module from the system, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supply from the source, and then unplug the cable from the power supply.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

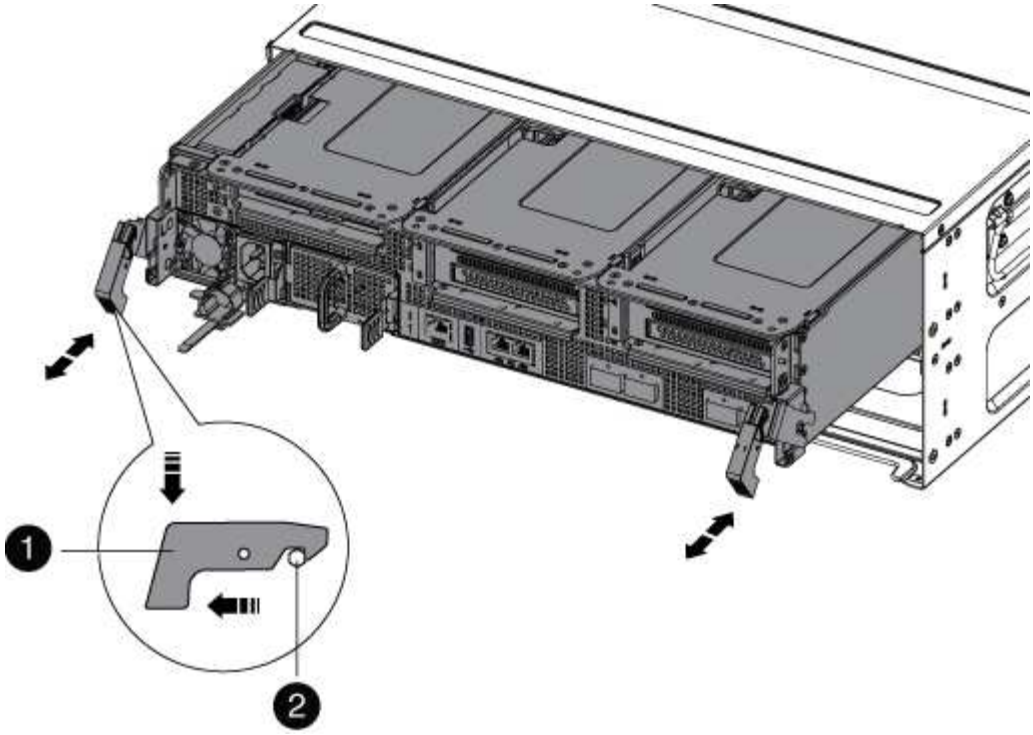

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Set the controller module aside in a safe place.

Step 3: Replace the NVRAM battery

To replace the NVRAM battery, you must remove the failed NVRAM battery from the controller module and install the replacement NVRAM battery into the controller module.

1. If you are not already grounded, properly ground yourself.
2. Locate the NVRAM battery on the left side of the riser module, Riser 1.

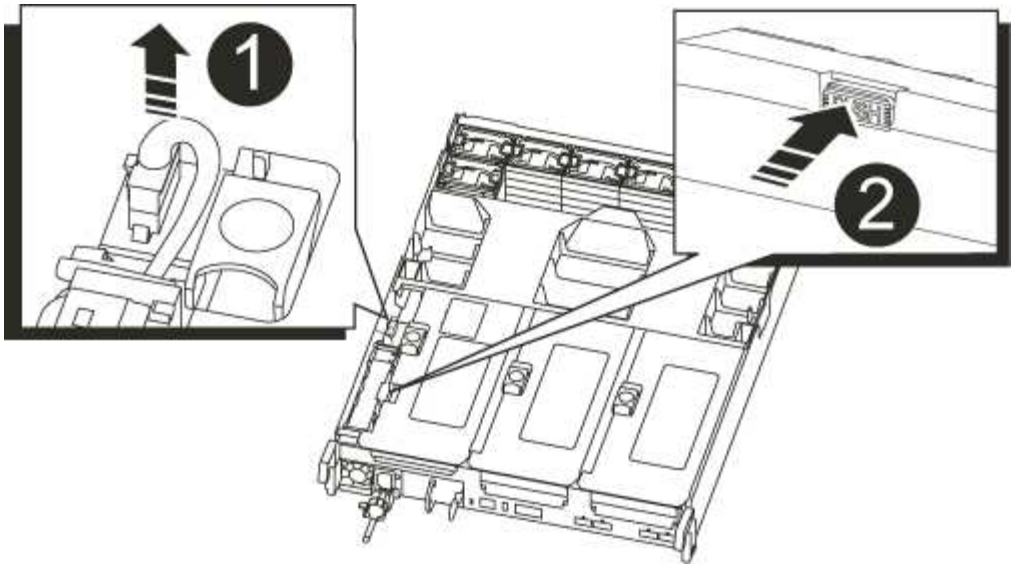

1	NVRAM battery plug
2	Blue NVRAM battery locking tab

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Push the blue locking tab on the battery holder, so that the latch releases from the holder.
5. Slide the battery down the riser bracket, lift the battery out of the controller, and then set it aside.
6. Slide the replacement battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and locks into place.
7. Plug the battery plug into the riser socket and make sure that the plug locks into place.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

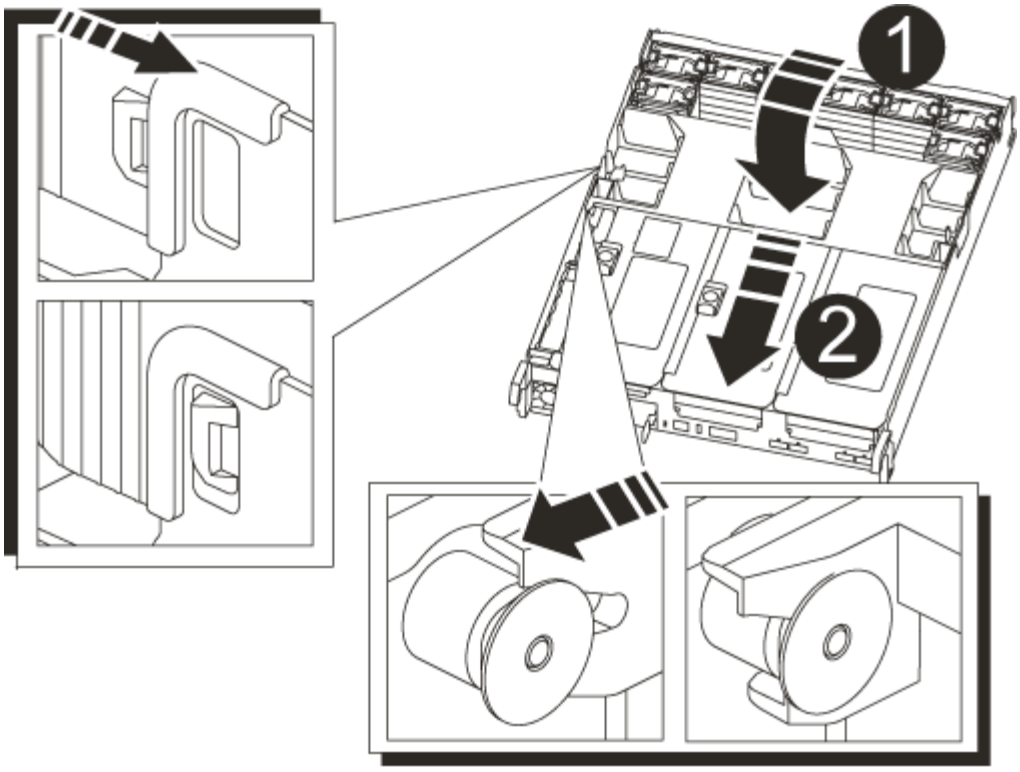

1	Locking tabs
2	Slide plunger

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   d. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect

the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

6. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs or onboard ports, convert these ports to 10 GbE connections by using the nicadmin convert command from Maintenance mode.

   Be sure to exit Maintenance mode after completing the conversion.

7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVRAM module and NVRAM DIMMs - AFF A700s

To replace a failed NVRAM card, you must remove the NVRAM riser, Riser 1, from the controller module, remove the failed card from the riser, install the new NVRAM card in the riser, and then reinstall the riser in the controller module. Because the system ID is derived from the NVRAM card, if replacing the module, disks belonging to the system are reassigned to the new system ID.

Before you begin

- All disk shelves must be working properly.
- If your system is in an HA pair, the partner controller must be able to take over the controller associated with the NVRAM module that is being replaced.
- This procedure uses the following terminology:
  - The *impaired* controller is the controller on which you are performing maintenance.
  - The *healthy* controller is the HA partner of the impaired controller.
- This procedure includes steps for automatically or manually reassigning disks to the controller module associated with the new NVRAM module. You must reassign the disks when directed to in the procedure. Completing the disk reassignment before giveback can cause issues.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You cannot change any disks or disk shelves as part of this procedure.

Step 1: Shut down the impaired controller

Steps

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

   - If you are using NetApp Storage Encryption, you must have reset the MSID using the instructions in the "Returning SEDs to unprotected mode" section of the *ONTAP 9 NetApp Encryption Power Guide*.

     ONTAP 9 NetApp Encryption Power Guide

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Unplug the controller module power supply from the source, and then unplug the cable from the power supply.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

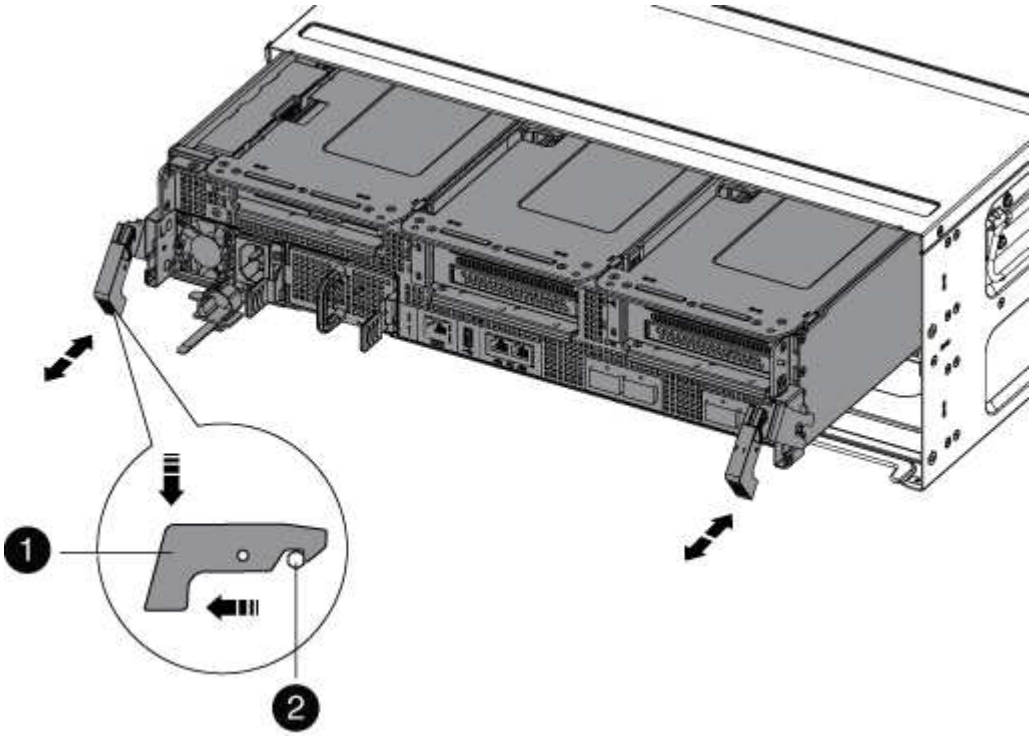

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

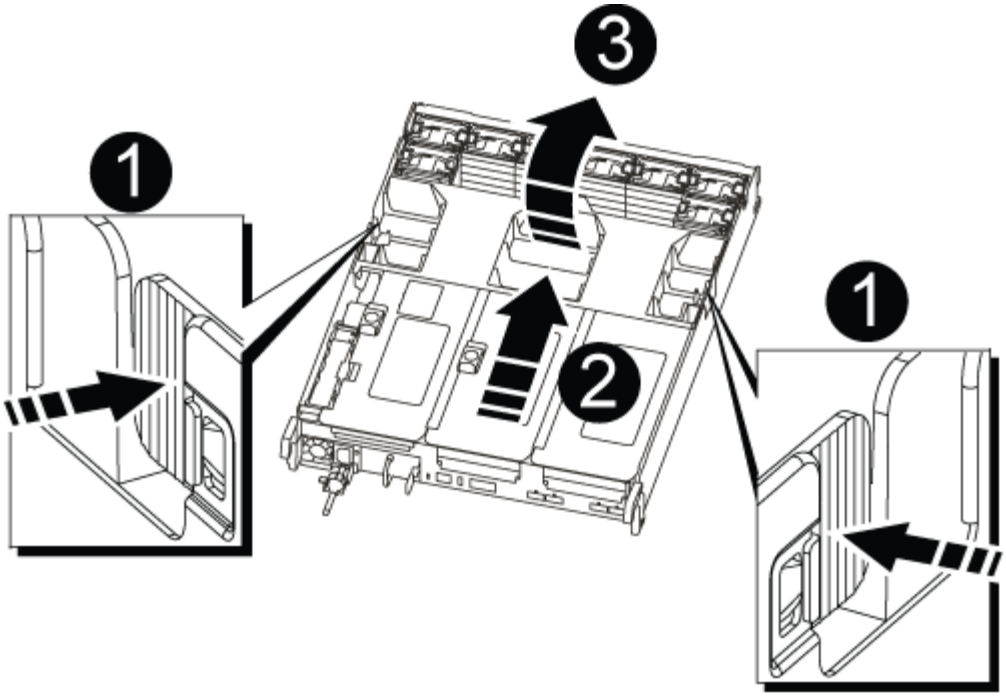

1	Air duct locking tabs
2	Risers
3	Air duct

Step 3: Remove the NVRAM card

Replacing the NVRAM consist of removing the NVRAM riser, Riser 1, from the controller module, disconnecting the NVRAM battery from the NVRAM card, removing the failed NVRAM card and installing the replacement NVRAM card, and then reinstalling the NVRAM riser back into the controller module.

1. If you are not already grounded, properly ground yourself.
2. Remove the NVRAM riser, Riser 1, from the controller module:
   a. Rotate the riser locking latch on the left side of the riser up and toward the fans.

      The NVRAM riser raises up slightly from the controller module.

   b. Lift the NVRAM riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser straight up out of the controller module, and then place it on a stable, flat surface so that you can access the NVRAM card.

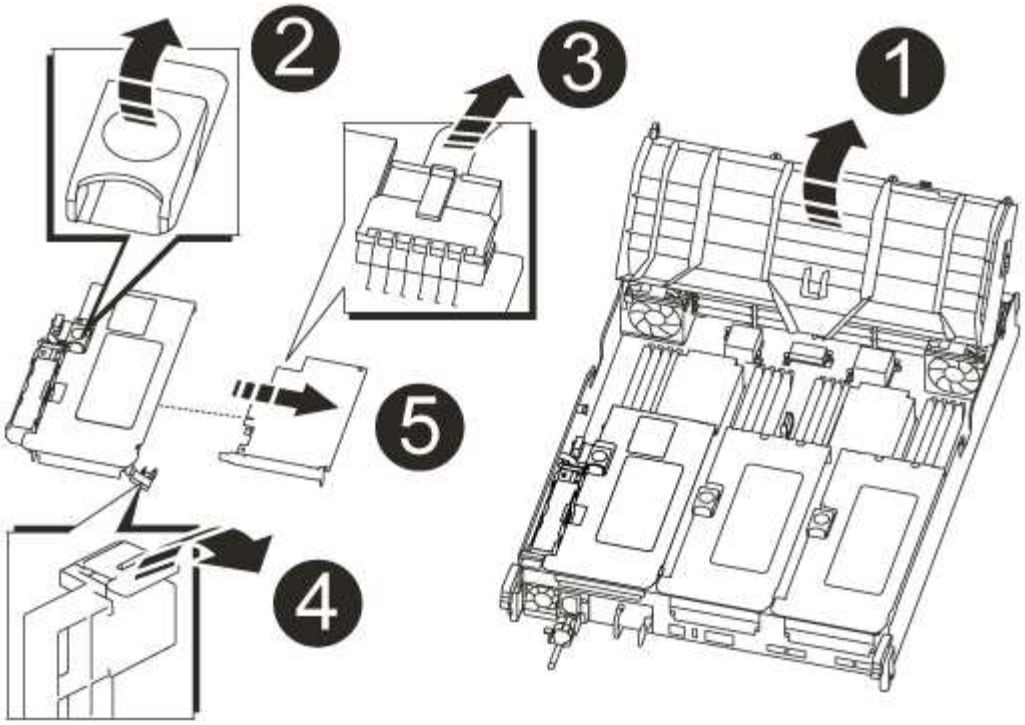

1	Air duct
2	Riser 1 locking latch
3	NVRAM battery cable plug connecting to the NVRAM card
4	Card locking bracket
5	NVRAM card

3. Remove the NVRAM card from the riser module:
   a. Turn the riser module so that you can access the NVRAM card.
   b. Unplug the NVRAM battery cable that is attached to the NVRAM card.
   c. Press the locking bracket on the side of the NVRAM riser, and then rotate it to the open position.
   d. Remove the NVRAM card from the riser module.
4. Install the NVRAM card into the NVRAM riser:
   a. Align the card with the card guide on the riser module and the card socket in the riser.
   b. Slide the card squarely into the card socket.

      Make sure that the card is completely and squarely seated into the riser socket.

   c. Connect the battery cable to the socket on the NVRAM card.

d. Swing the locking latch into the locked position and make sure that it locks in place.

5. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

   d. Reinsert any SFP modules that were removed from the PCIe cards.

Step 4: Reinstall the controller module and booting the system

After you replace a FRU in the controller module, you must reinstall the controller module and reboot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. If you are not already grounded, properly ground yourself.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   d. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   e. Interrupt the boot process by pressing `Ctrl-C` when you see `Press Ctrl-C for Boot Menu`.
   f. Select the option to boot to Maintenance mode from the displayed menu.

Step 5: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

   ```
   node1> `storage failover show`
                                       Takeover
   Node              Partner           Possible     State Description
   ------------      ------------      --------
   -------------------------------------
   node1             node2             false        System ID changed on
   partner (Old:
                                                     151759755, New:
   151759706), In takeover
   node2             node1             -            Waiting for giveback
   (HA mailboxes)
   ```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

The *replacement* controller takes back its storage and completes booting.

If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration Guide for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

6. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

7. Verify that the expected volumes are present for each controller: `vol show -node node-name`
8. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe card - AFF A700s

To replace a PCIe card, you must disconnect the cables from the cards in the riser, remove the riser, replace the riser, and then recable the cards in that riser.

- You can use this procedure with all versions of ONTAP supported by your system

- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Unplug the controller module power supply from the source, and then unplug the cable from the power

supply.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

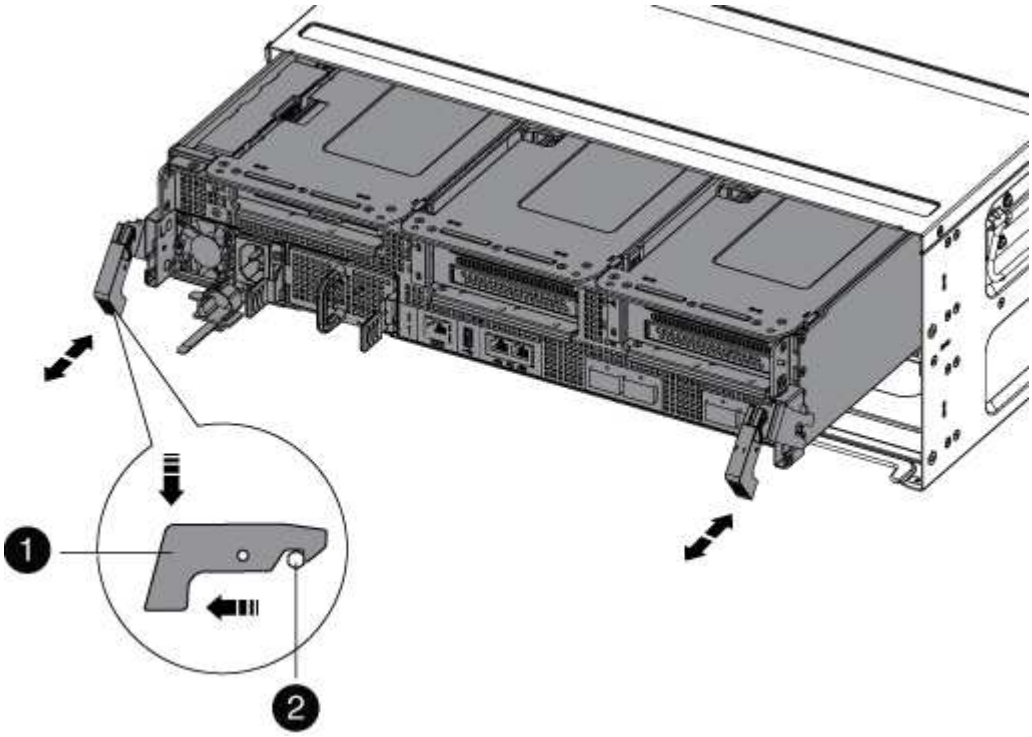

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface, and then open the air duct:

   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

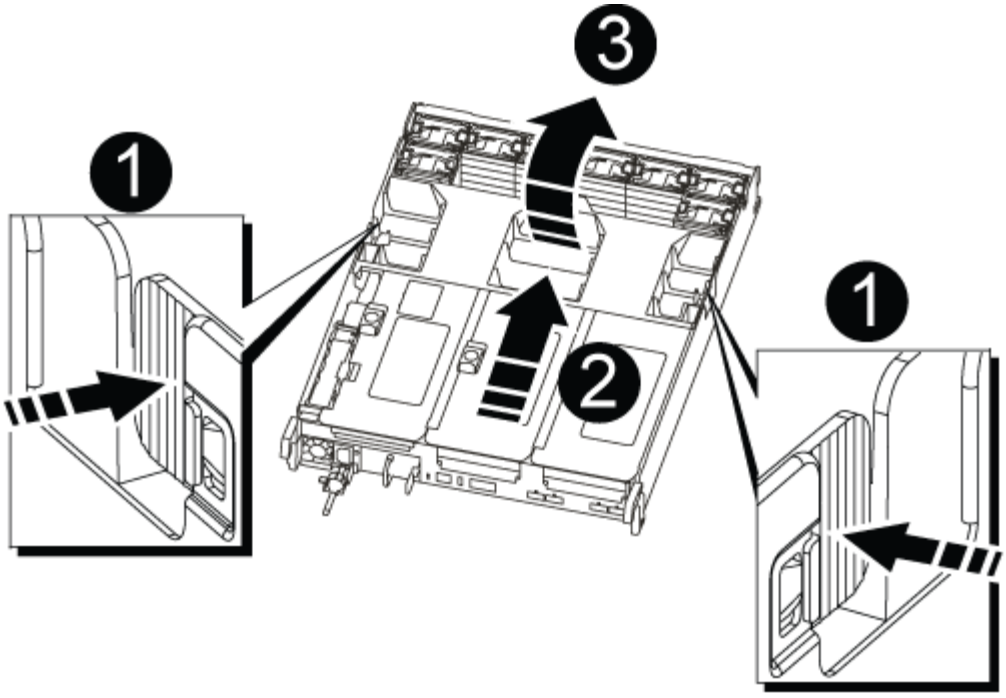

1	Air duct locking tabs
2	Risers
3	Air duct

Step 3: Replace a PCIe card

To replace a PCIe card, you must remove the cabling and any SFPs from the ports on the PCIe cards in the target riser, remove the riser from the controller module, remove and replace the PCIe card, reinstall the riser, and recable it.

1. If you are not already grounded, properly ground yourself.
2. Remove the PCIe riser from the controller module:
   a. Remove any SFP modules that might be in the PCIe cards.
   b. Rotate the module locking latch on the left side of the riser up and toward the fan modules.

      The PCIe riser raises up slightly from the controller module.

   c. Lift the PCIe riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

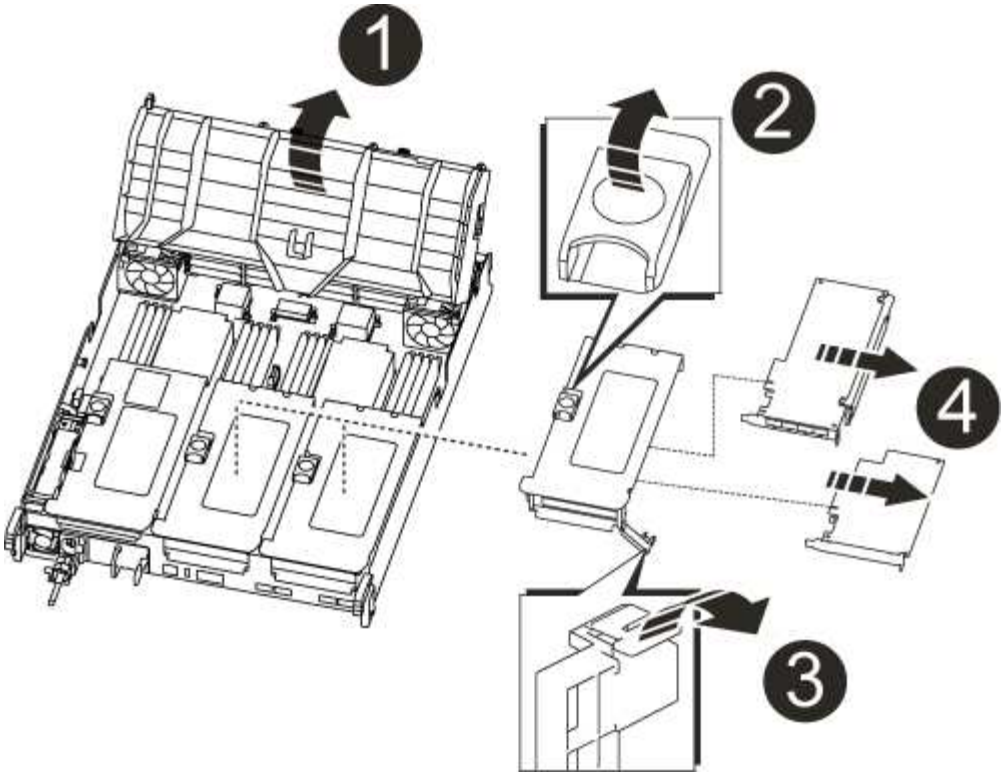

1	Air duct
2	Riser locking latch
3	Card locking bracket
4	Riser 2 (middle riser) and PCI cards in riser slots 2 and 3.

3. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. Remove the PCIe card from the riser.
4. Install the PCIe card into the same slot in PCIe riser:
   a. Align the card with the card guide on the riser and the card socket in the riser, and then slide it squarely into the socket in the riser.

      Make sure that the card is completely and squarely seated into the riser socket.

   b. Swing the locking latch into place until it clicks into the locked position.
5. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller

module.

c. Swing the locking latch down and click it into the locked position.

   When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, close the air duct:
   a. Swing the air duct all the way down to the controller module.
   b. Slide the air duct toward the risers until the locking tabs click into place.
   c. Inspect the air duct to make sure that it is properly seated and locked into place.

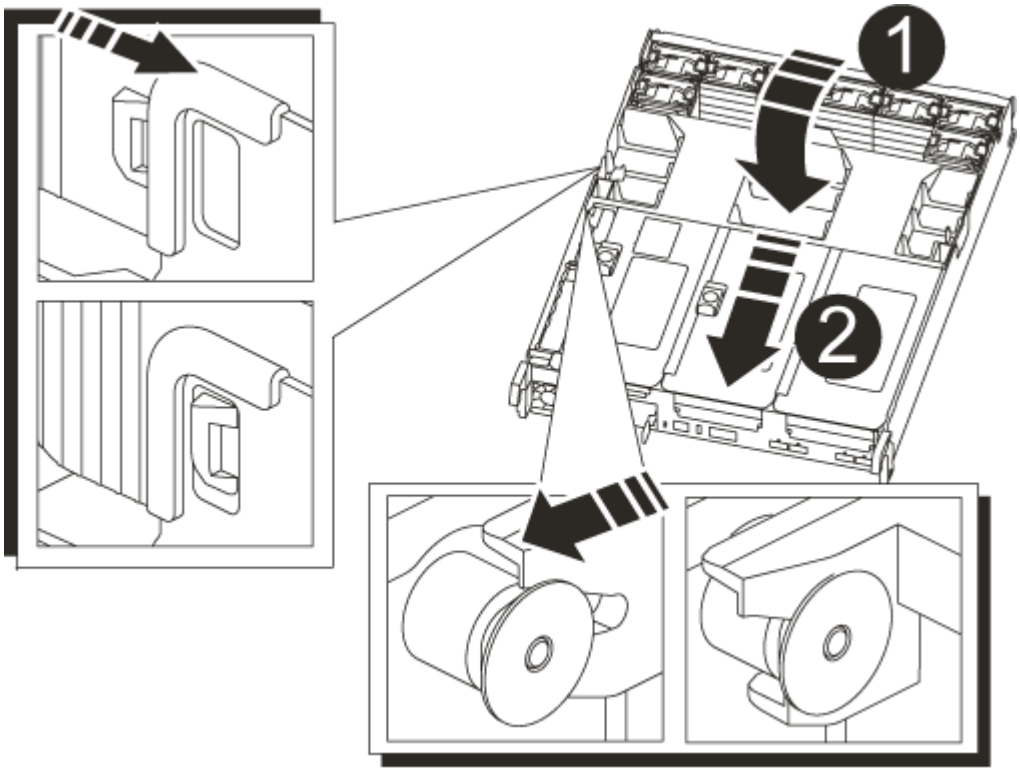

1	Locking tabs
2	Slide plunger

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

4. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   d. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

6. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs or onboard ports, convert these ports to 10 GbE connections by using the nicadmin convert command from Maintenance mode.

   Be sure to exit Maintenance mode after completing the conversion.

7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - AFF A700s

Swapping out a power supply involved disconnecting the target power supply (PSU) from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- The number of power supplies in the system depends on the model.
- Power supplies are auto-ranging.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
3. Disconnect the power supply:
   a. Open the power cable retainer, and then unplug the power cable from the power supply.
   b. Unplug the power cable from the power source.
4. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

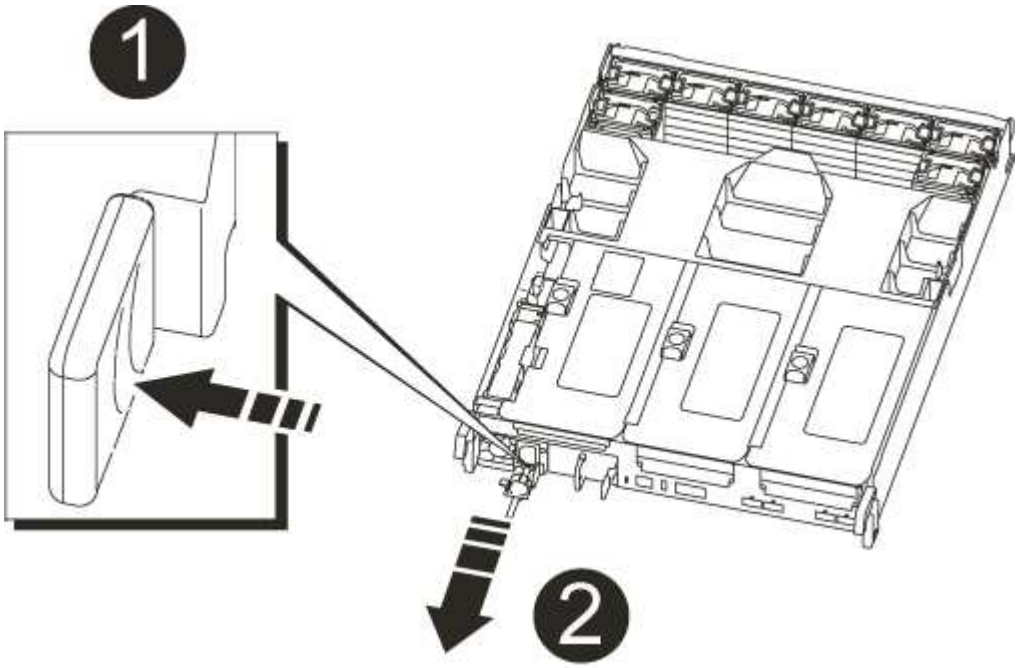

1	Blue power supply locking tab
2	Power supply

5. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into

place.

The power supplies will only properly engage with the internal connector and lock in place one way.

> To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

6. Close the cam handle by swinging it down as far as it will go.
7. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF A700s

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`
2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Unplug the controller module power supply from the source, and then unplug the cable from the power supply.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

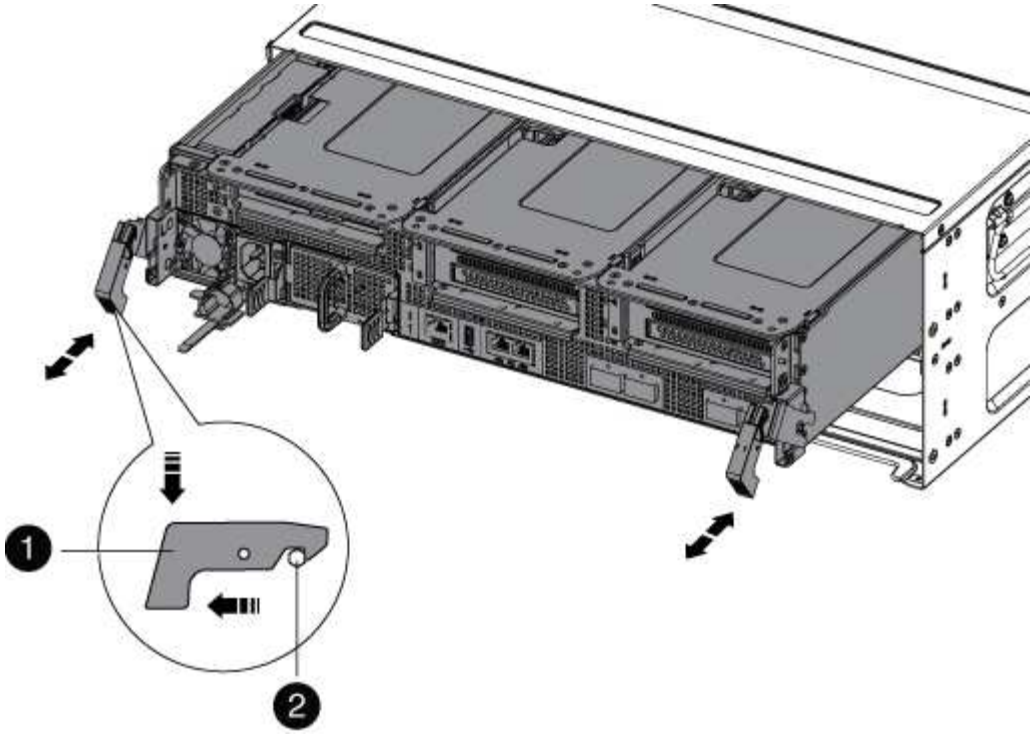

1	Locking latch
2	Locking pin

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

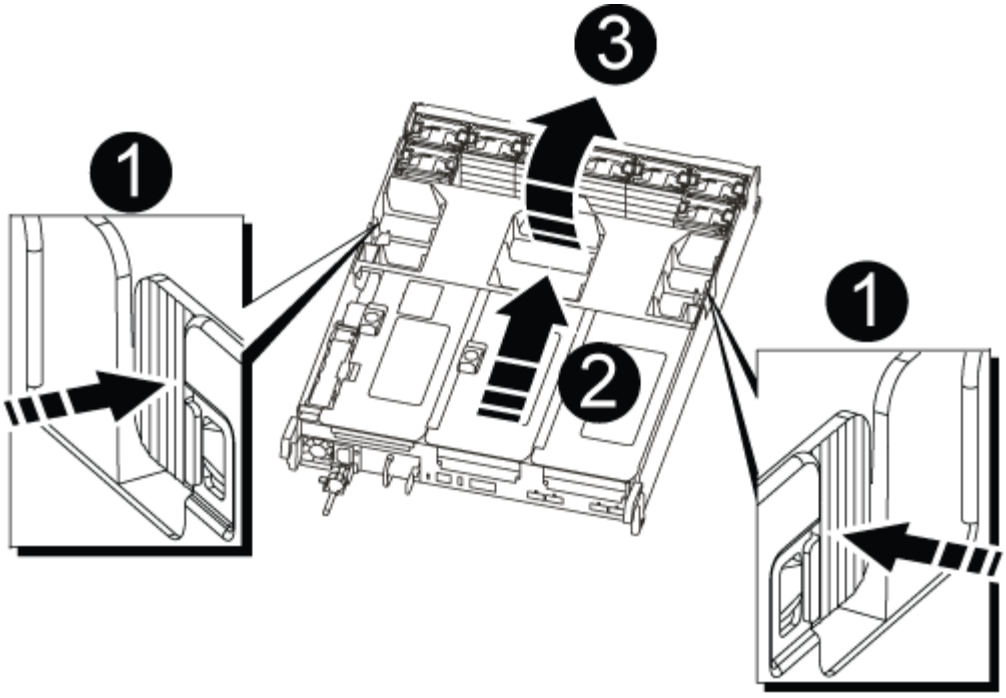

1	Air duct locking tabs
2	Risers
3	Air duct

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

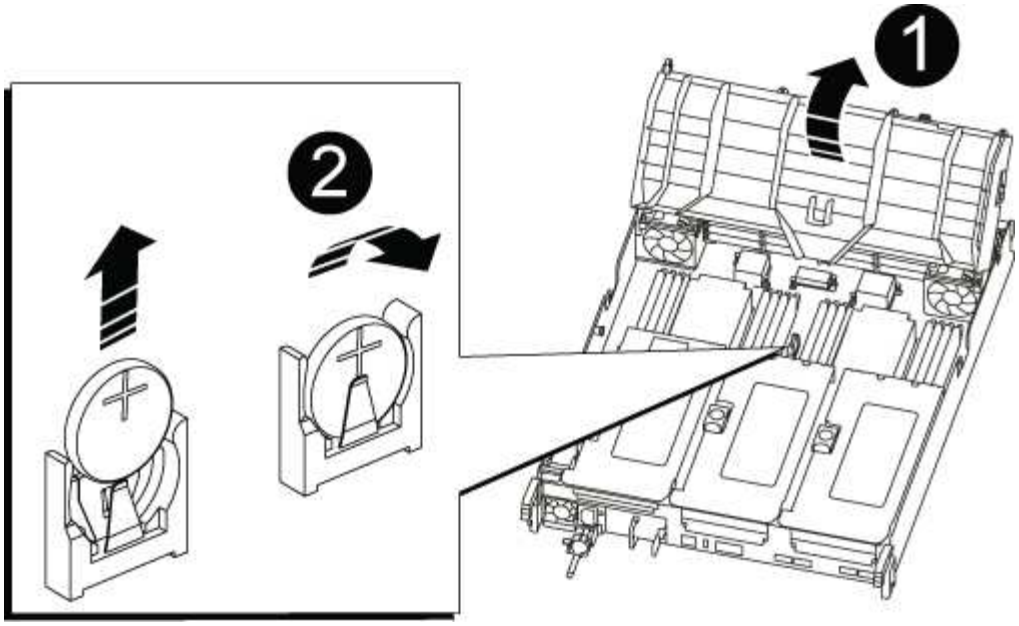

1	Air duct
2	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber

optic cables.

4. Complete the reinstallation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Halt the controller at the LOADER prompt.
5. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

AFF C190 systems

Install and setup

Start here: Choose your installation and setup experience

You can choose from different content formats to guide you through installing and setting up your new storage system.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

Quick steps - AFF C190

Learn how to install your C190 system from racking and cabling, through initial system bring-up. Use the AFF C190 Installation and Setup Instructions if you are familiar with installing NetApp systems.

Video steps - AFF C190

The following video shows how to install and cable your system.

Animation - Install and Setup of an AFF C190

Detailed steps - AFF C190

This section gives detailed step-by-step instructions for installing a AFF C190 system.

Step 1: Prepare for installation

To install your AFF C190 system, create an account and register the system. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

- Make sure you have access to NetApp Hardware Universe (HWU) for information about site requirements as well as additional information on your configured system.
- Make sure you have access to the Release Notes for your version of ONTAP for more information about this system.
- Make sure you have the following items at your site:
  - Rack space for the storage system
  - Phillips #2 screwdriver
  - Additional networking cables to connect your system to your network switch and laptop or console with a Web browser
  - A laptop or console with an RJ-45 connection and access to a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register (NetApp Product Registration) your system.
4. Download and install NetApp Downloads: Config Advisor on your laptop.
5. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
10 GbE cable (order dependent)	X6566B-05-R6 (112-00297), 0.5m X6566B-2-R6 (112-00299), 2m		Cluster interconnect network
	X6566B-2-R6 (112-00299), 2m X6566B-3-R6 (112-00300), 3m X6566B-5-R6 (112-00301), 5m		Data
Optical network cables (order dependent)	X6553-R6 (112-00188), 2m X6536-R6 (112-00090), 5m X6554-R6(112-00189), 15m	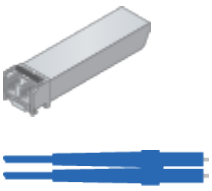	SFP + FC host network
Cat 6, RJ-45 (order dependent)	X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Ethernet host and management network
Micro-USB console cable	Not applicable		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	Not applicable		Powering up the system

6. Download and complete the Cluster Configuration Worksheet.

Step 2: Install the hardware

Install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

You need to be aware of the safety concerns associated with the weight of the system.

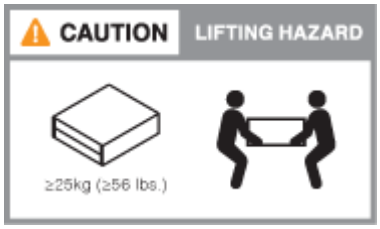

3. Attach cable management devices (as shown).

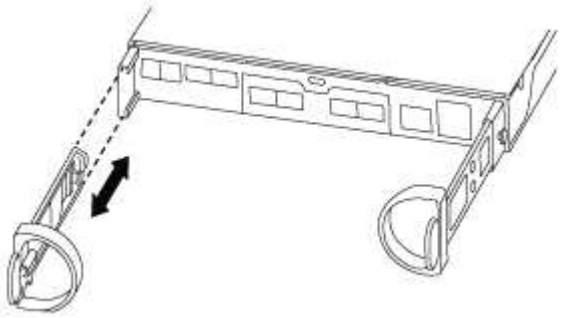

4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

Cable the controllers to your network by using the two-node switchless cluster method or the cluster interconnect network method.

The following table identifies the cable type with the call out number and cable color in the illustrations for both two-node switchless cluster network cabling and switched cluster network cabling.

Cabling	Connection type
1	Cluster interconnect
2	Controllers to host data network switches
3	Controllers to management network switch

Option 1: Two-node switchless cluster

Learn how to cable a two-node switchless cluster.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

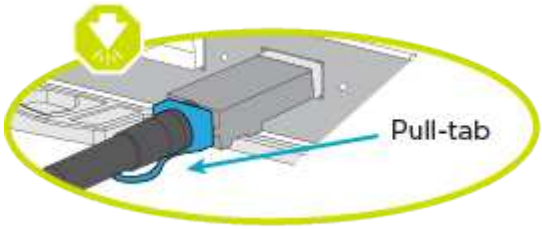

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

If connecting to an optical switch, insert the SFP into the controller port before cabling to the port.

About this task

Refer to the following cabling illustrations when cabling between the controllers and the switches.

UTA2 data network configurations

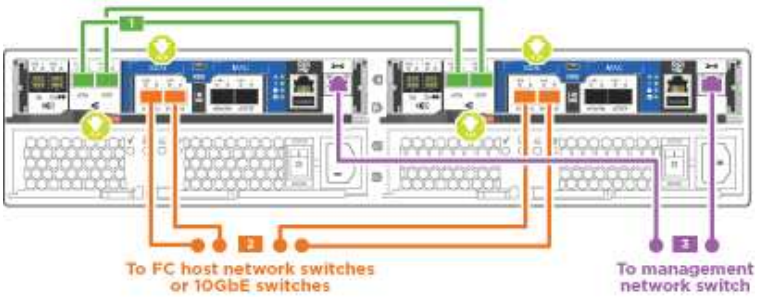

Ethernet network configurations

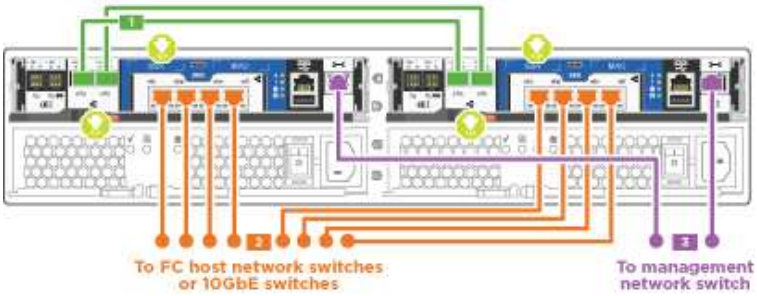

Perform the following steps on each controller module.

Steps

1. Cable the cluster interconnect ports e0a to e0a and e0b to e0b with the cluster interconnect cable.

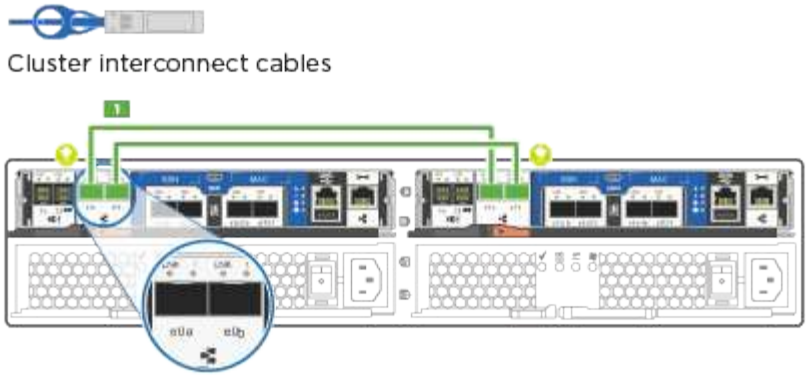

2. Cable the controllers to either a UTA2 data network or an Ethernet network.

UTA2 data network configurations

Use one of the following cable types to cable the e0c/0c and e0d/0d or e0e/0e and e0f/0f data ports to your host network.

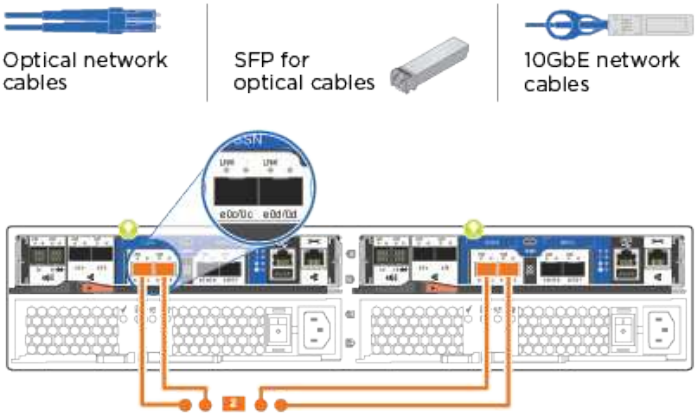

Ethernet network configurations

Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network.
in the following illustration.

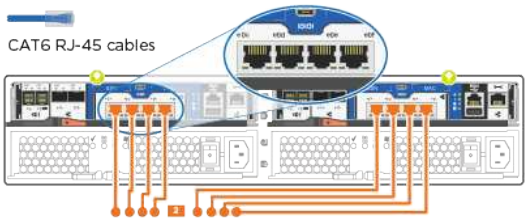

3. Cable the e0M ports to the management network switches with the RJ45 cables.

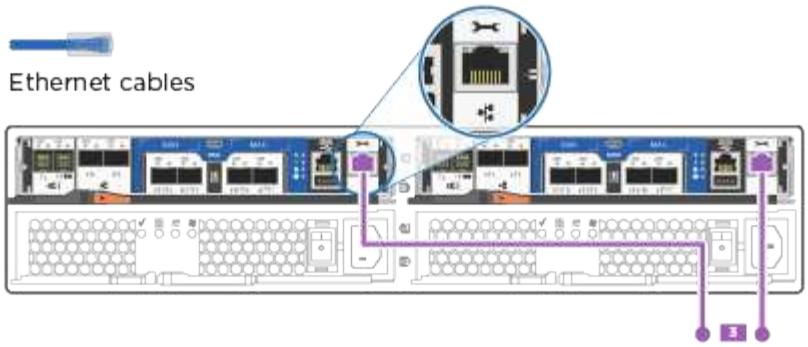

DO NOT plug in the power cords at this point.

Option 2: Switched cluster

Learn how to cable a switched cluster.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

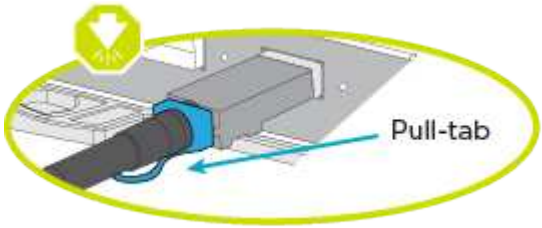

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

If connecting to an optical switch, insert the SFP into the controller port before cabling to the port.

About this task

Refer to the following cabling illustrations when cabling between the controllers and the switches.

Unified network configurations

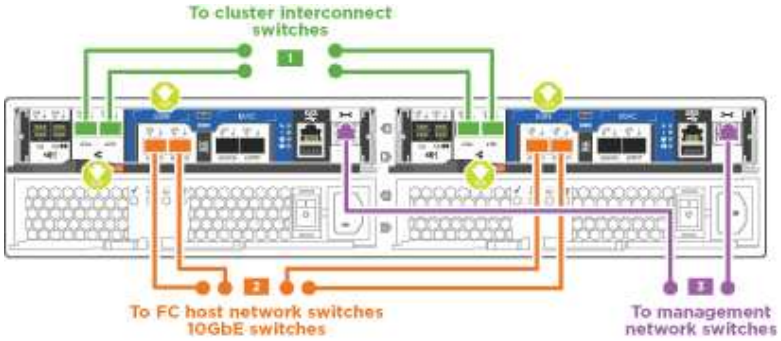

Ethernet network configurations

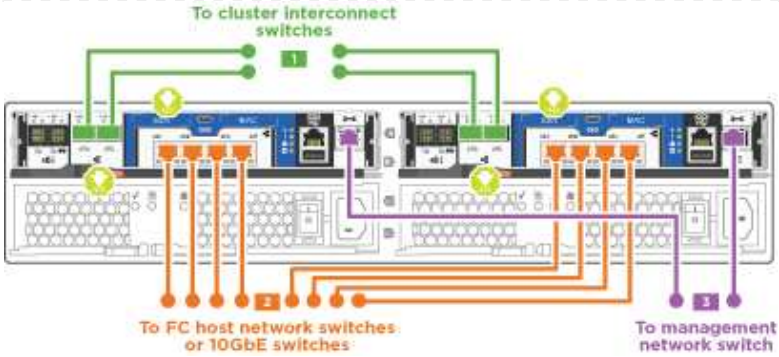

Perform the following steps on each controller module.

Steps

1. Cable e0a and e0b to the cluster interconnect switches with the cluster interconnect cable.

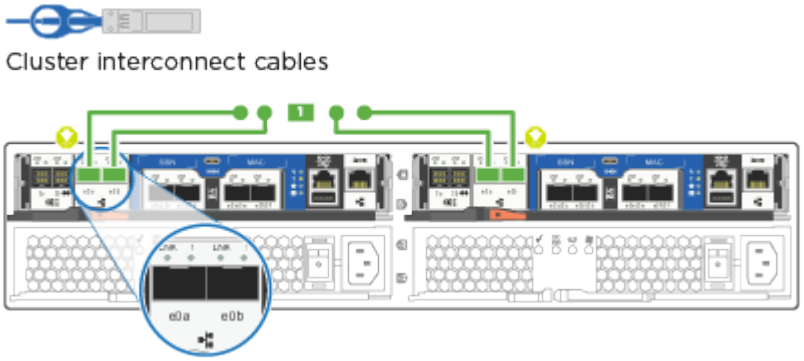

2. Cable the controllers to either a UTA2 data network or an Ethernet network.

UTA2 data network configurations

Use one of the following cable types to cable the e0c/0c and e0d/0d or e0e/0e and e0f/0f data ports to your host network.

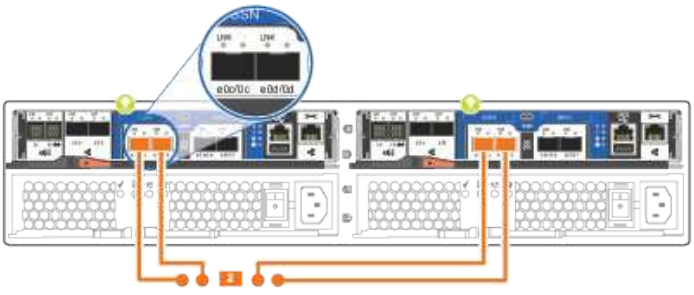

Ethernet network configurations	Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network.

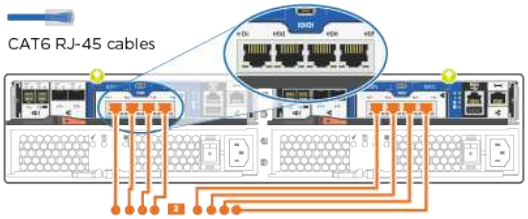

3. Cable the e0M ports to the management network switches with the RJ45 cables.

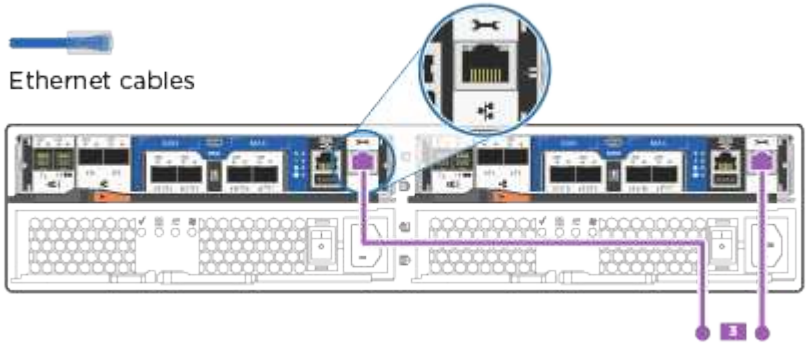

DO NOT plug in the power cords at this point.

Step 4: Complete system setup

Complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: If network discovery is enabled

Learn how to complete system setupIf you have network discovery enabled on your laptop.

Steps

1. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
2. Turn on the power switches to both nodes.

Initial booting may take up to eight minutes..

3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Connect your laptop to the Management switch:

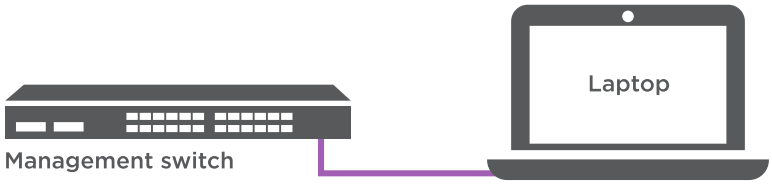

1. Select an ONTAP icon listed to discover:

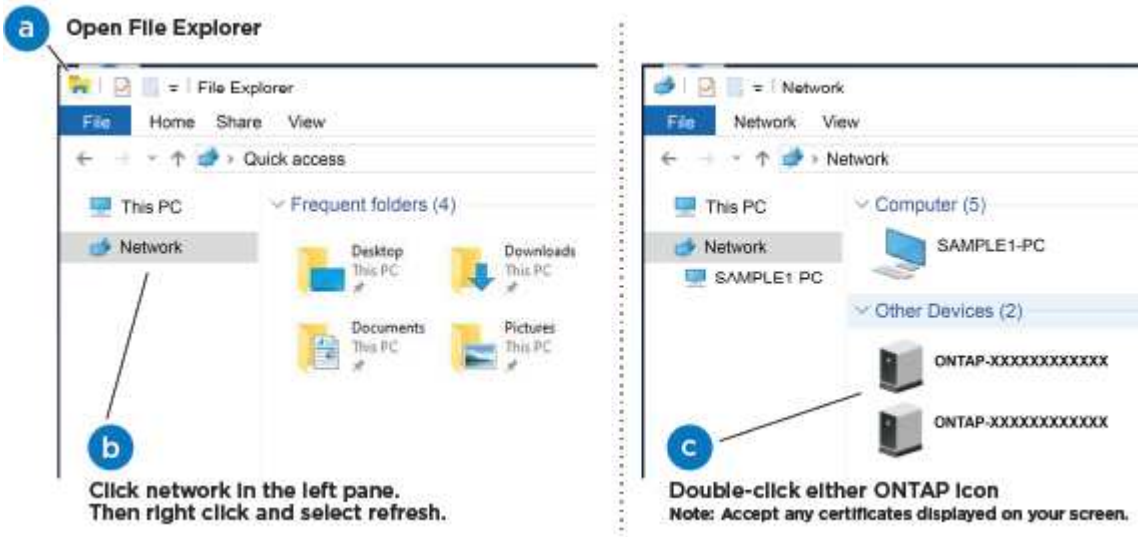

a. Open File Explorer.

b. Click **Network** in the left pane.

c. Right-click and select **refresh**.

d. Double-click either ONTAP icon and accept any certificates displayed on your screen.

XXXXX is the system serial number for the target node.

System Manager opens.

2. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
3. Verify the health of your system by running Config Advisor.
4. After you have completed the initial configuration, go to the ONTAP documentation site for information about configuring additional features in ONTAP.

The default port configuration for Unified configuration systems is CNA mode; if connecting to an FC host network, you have to modify the ports for FC mode.

Option 2: If network discovery is not enabled

Learn how to complete the system setup if network discovery is not enabled on your laptop.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

   See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your system.

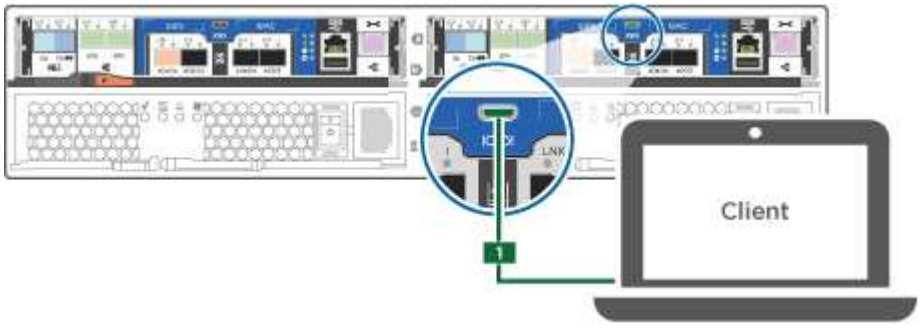

   c. Connect the laptop or console to the switch on the management subnet.

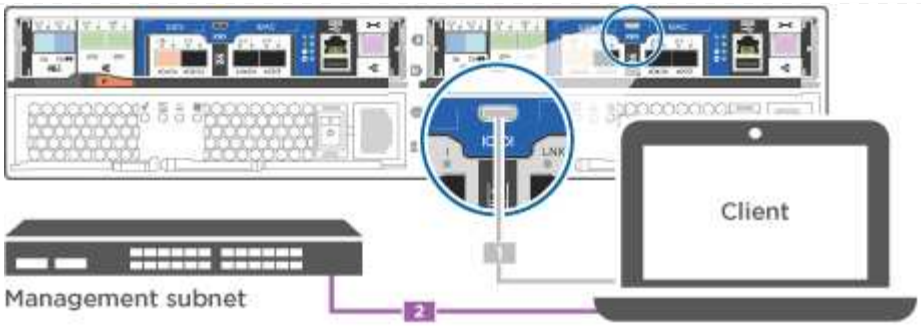

d. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Turn on the power switches to both nodes.

Initial booting may take up to eight minutes..

4. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

5. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

   The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the ONTAP Configuration Guide.
6. Verify the health of your system by running Config Advisor.

7. After you have completed the initial configuration, go to the ONTAP &ocumentation site for information about configuring additional features in ONTAP.

> The default port configuration for Unified configuration systems is CNA mode; if connecting to an FC host network, you have to modify the ports for FC mode.

Maintain

Maintain AFF C190 hardware

For the AFF C190 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

Fan

The fan cools the controller.

NVMEM battery

A battery is included with the controller and provides the backup power if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real-time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF C190

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the var file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the var file system.
  - For disruptive replacement, you do not need a network connection to restore the var file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
  - The *impaired* controller is the controller on which you are performing maintenance.
  - The *healthy* controller is the HA partner of the impaired controller.

Check encryption key support and status - AFF C190

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key

Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller - AFF C190

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller displays…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Replace the boot media - AFF C190

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller

To access components inside the controller module, you must first remove the controller module from the system, and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

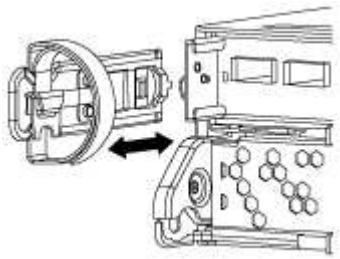

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

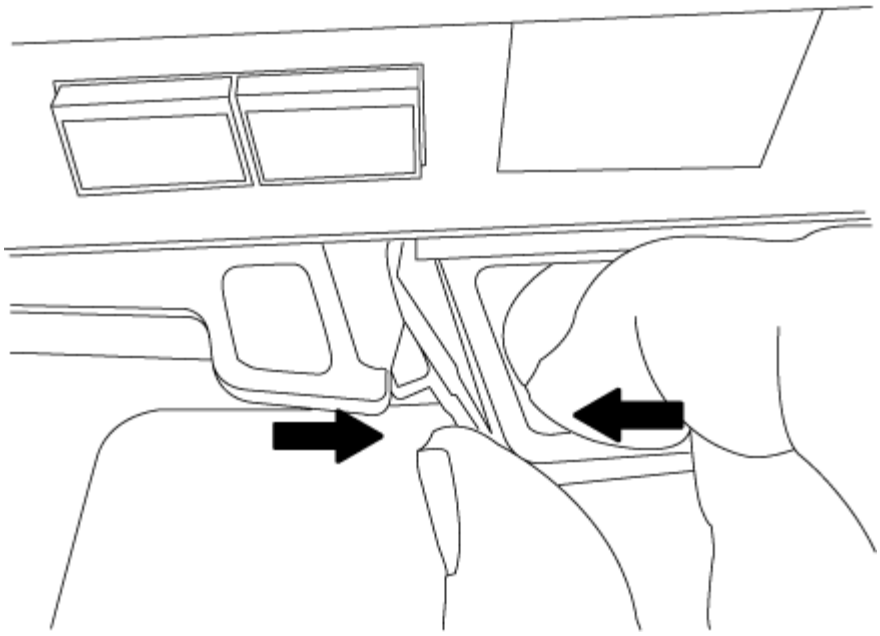

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

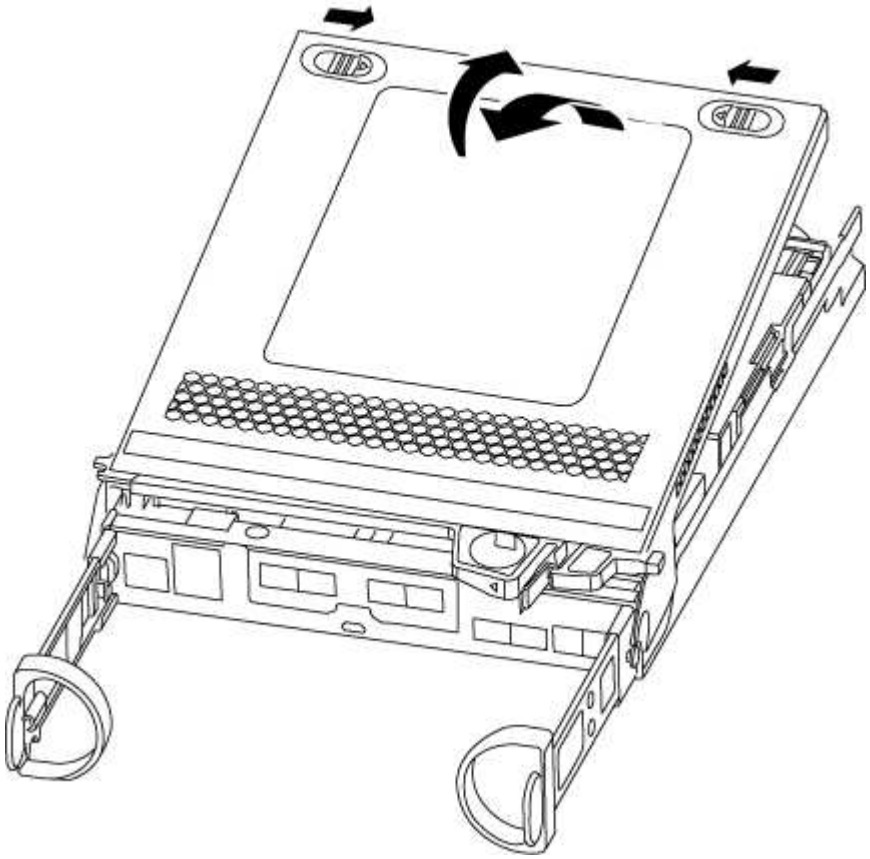

Step 2: Replace the boot media

You must locate the boot media in the controller module, and then follow the directions to replace it.

1. Locate the boot media using the following illustration or the FRU map on the controller module:
2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   > Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.
6. Close the controller module cover.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the `var` file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the **Downloads** section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

   The controller begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see `Starting AUTOBOOT press Ctrl-C to abort….`

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then `halt` the controller to boot to LOADER.

6. Boot the recovery image:

   `boot_recovery` *`ontap_image_name`*`.tgz`

   > If the `image.tgz` file is named something other than image.tgz, such as `boot_recovery 9_4.tgz`, you need to include the different file name in the `boot_recovery` command.

   The system boots to the boot menu and prompts you for the boot image name.

7. Enter the boot image name that is on the USB flash drive:

   ***`image_name`*`.tgz`**

   After `image_name.tgz` is installed, the system prompts you to restore the backup configuration (the `var` file system) from the healthy controller.

8. Restore the `var` file system:

If your system has…	Then…
A network connection	a. Press **`y`** when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: **`set -privilege advanced`** c. Run the restore backup command: **`system node restore-backup -node local -target -address`** ***`impaired_node_IP_address`*** d. Return the controller to admin level: **`set -privilege admin`** e. Press **`y`** when prompted to use the restored configuration. f. Press **`y`** when prompted to reboot the controller.
No network connection	a. Press **`n`** when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press **`y`**.

9. Verify that the environmental variables are set as expected.
   a. Take the controller to the LOADER prompt.

      From the ONTAP prompt, you can issue the command `system node halt -skip-lif -migration-before-shutdown true -ignore-quorum-warnings true -inhibit -takeover true`.

   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment_variable_name changed_value`* command.
   d. Save your changes using the `saveenv` command.
   e. Reboot the controller.
10. The next step depends on your system configuration:

If your system is in…	Then…
A stand-alone configuration	You can begin using your system after the controller reboots.
An HA pair	After the impaired controller is displaying the `Waiting for Giveback…` message, perform a giveback from the healthy controller: a. Perform a giveback from the healthy controller: **`storage failover giveback -ofnode`** ***`partner_node_name`*** This initiates the process of returning ownership of the impaired controller's aggregates and volumes from the healthy controller back to the impaired controller. If the giveback is vetoed, you can consider overriding the vetoes. HA pair management b. Monitor the progress of the giveback operation by using the `storage failover show`-giveback` command. c. After the giveback operation is complete, confirm that the HA pair is healthy and that takeover is possible by using the `storage failover show` command. d. Restore automatic giveback if you disabled it by using the `storage failover modify` command.

Boot the recovery image - AFF C190

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive:

   `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the `var` file system:

If your system has…	Then…
A network connection	a. Press **`y`** when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: **`set -privilege advanced`** c. Run the restore backup command: **`system node restore-backup -node local -target -address impaired_node_IP_address`** d. Return the controller to admin level: **`set -privilege admin`** e. Press **`y`** when prompted to use the restored configuration. f. Press **`y`** when prompted to reboot the controller.
No network connection	a. Press **`n`** when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press **`y`**.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment_variable_name changed_value`* command.
   d. Save your changes using the `saveenv` command.
5. The next depends on your system configuration:
   - If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed

- If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.
8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore encryption - AFF C190

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
**********************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
**********************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - AFF C190

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - AFF C190

To replace the chassis, you must move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving all drives and controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF C190

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - AFF C190

Move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Move the power supply

Moving out a power supply when replacing a chassis involves turning off, disconnecting, and removing the power supply from the old chassis and installing and connecting it on the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.
4. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

Step 2: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

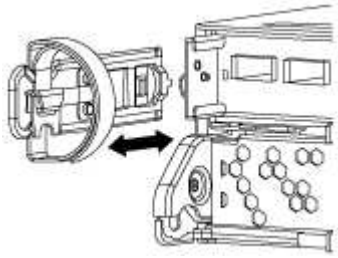

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

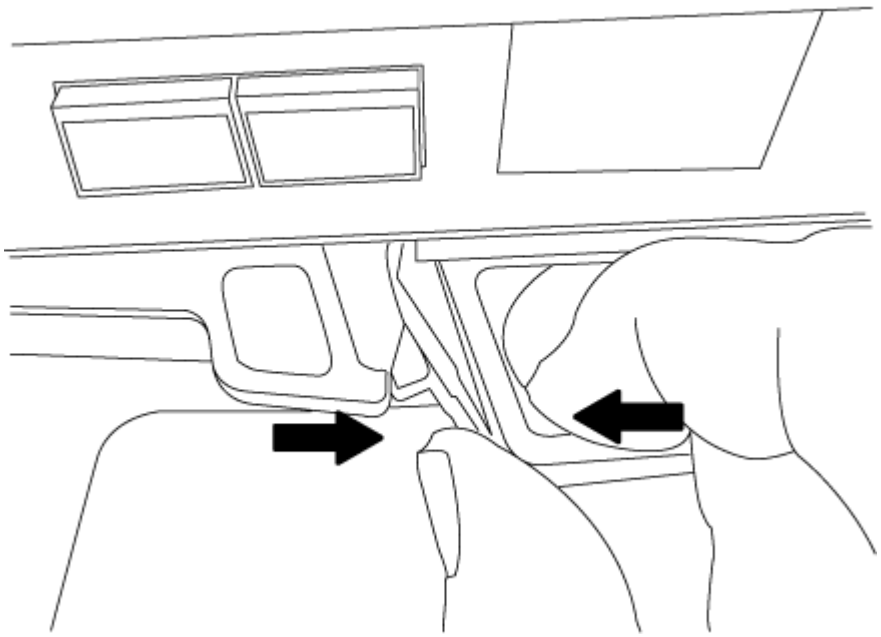

5. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 3: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.
5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.
6. Repeat the process for the remaining drives in the system.

Step 4: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller module

After you install the controller module and any other components into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Repeat the preceding steps if there is a second controller to install in the new chassis.
4. Complete the installation of the controller module
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Repeat the preceding steps for the second controller module in the new chassis.
5. Connect the power supplies to different power sources, and then turn them on.
6. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

      If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

   b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - AFF C190

You must verify the HA state of the chassis.

Step 1: Verify and setting the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis:

   `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

a. Set the HA state for the chassis:

ha-config modify chassis *HA-state*

The value for *HA-state* can be one of the following:

- ha
- non-ha

b. Confirm that the setting has changed:

ha-config show

3. If you have not already done so, recable the rest of your system.
4. Reboot the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - AFF C190

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- This procedure includes steps for automatically or manually reassigning drives to the *replacement* controller, depending on your system’s configuration.
- You should perform the drive reassignment as directed in the procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might

encounter during the replacement process.

Shut down the controller - AFF C190

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the controller module hardware - AFF C190

To replace the controller module, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove controller module

To replace the controller module, you must first remove the old controller module from the chassis.

1. If you are not already grounded, properly ground yourself.

2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

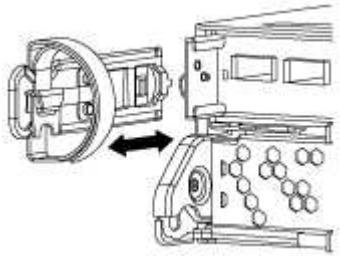

4. If you left the SFP modules in the system after removing the cables, move them to the new controller module.
5. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

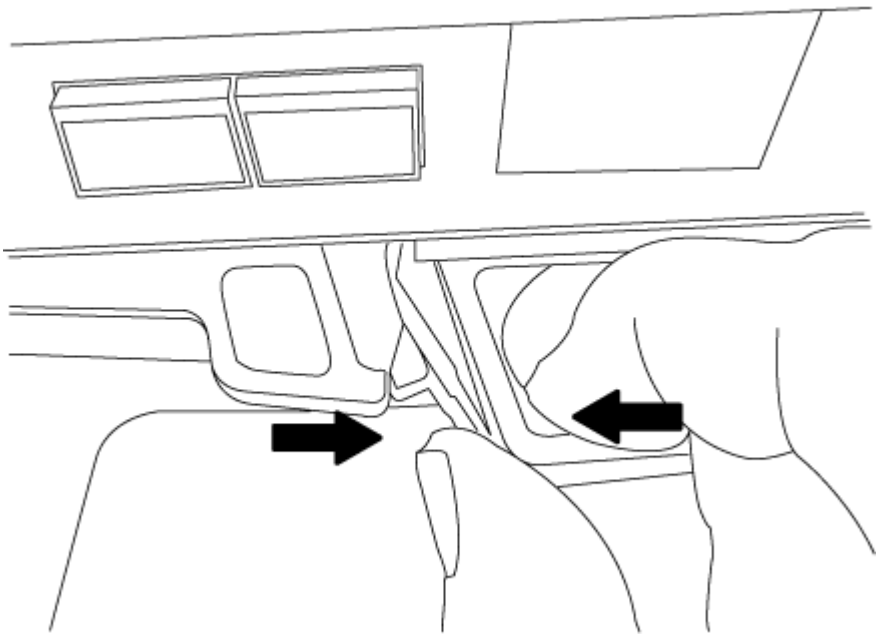

6. Turn the controller module over and place it on a flat, stable surface.
7. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

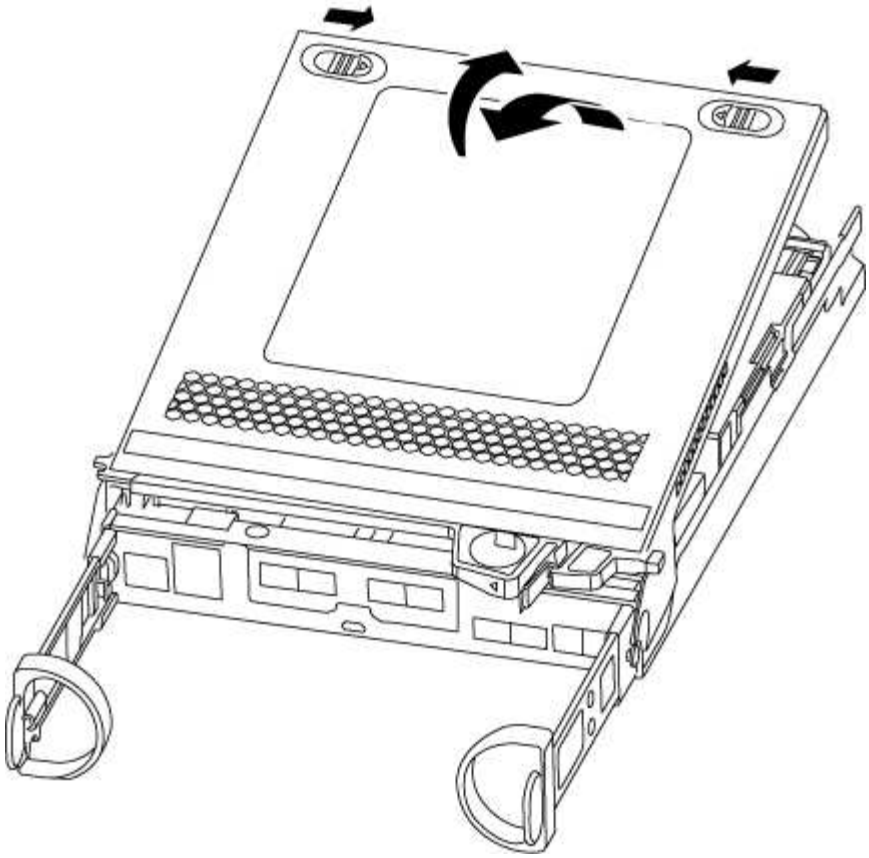

Step 2: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller module and insert it in the new controller module.

1. Locate the boot media using the following illustration or the FRU map on the controller module:

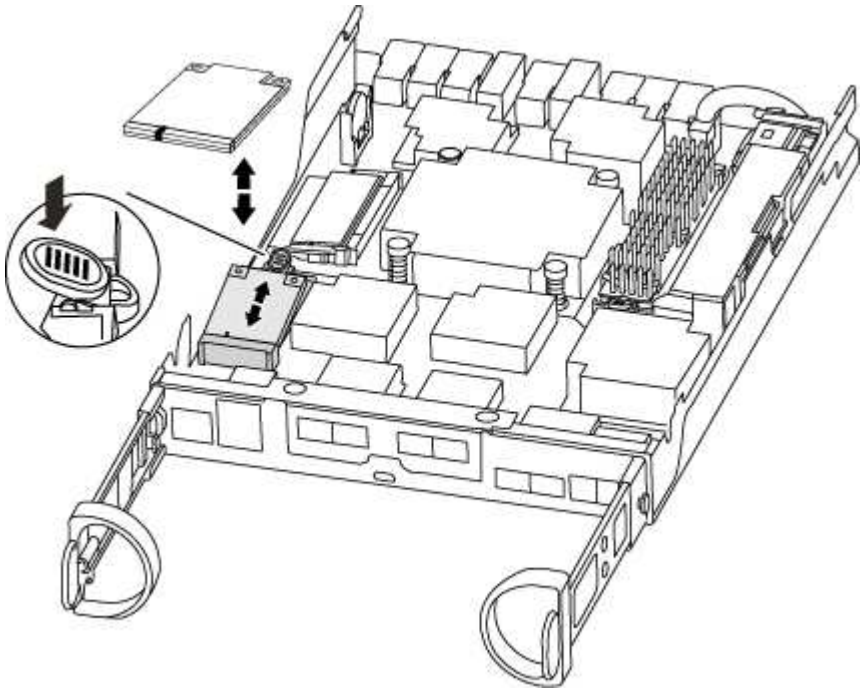

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the NVMEM battery

To move the NVMEM battery from the old controller module to the new controller module, you must perform a specific sequence of steps.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

   

   The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Locate the NVMEM battery in the controller module.

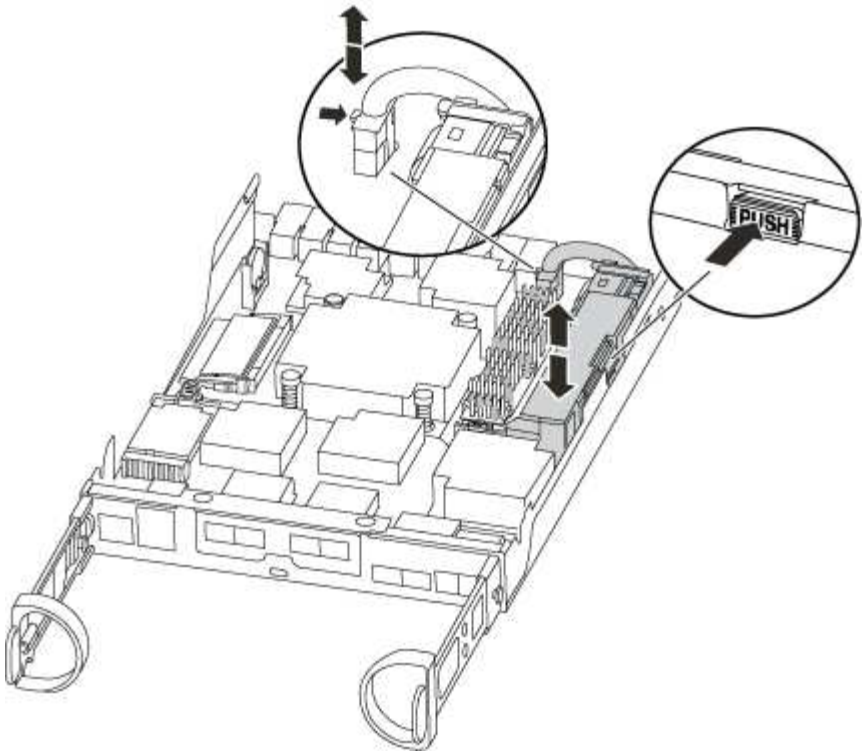

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Loop the battery cable around the cable channel on the side of the battery holder.
7. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
8. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.

Step 4: Move the DIMMs

To move the DIMMs, you must follow the directions to locate and move them from the old controller module into the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

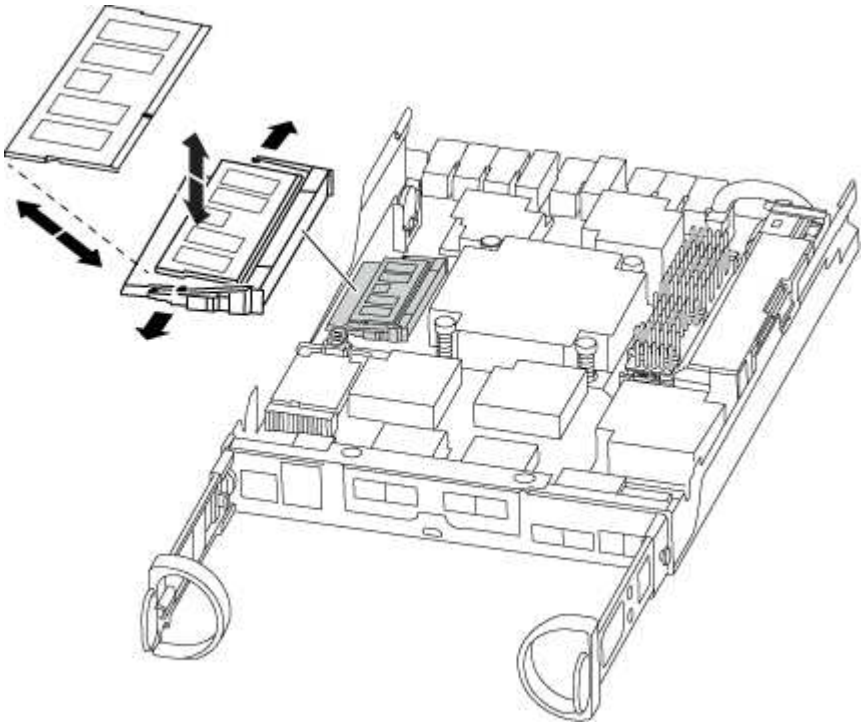

4. Repeat these steps to remove additional DIMMs as needed.
5. Verify that the NVMEM battery is not plugged into the new controller module.
6. Locate the slot where you are installing the DIMM.
7. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Repeat these steps for the remaining DIMMs.
9. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

   Make sure that the plug locks down onto the controller module.

Step 5: Install the controller module

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you have not already done so, replace the cover on the controller module.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the reinstallation of the controller module. The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller begins to boot as soon as it is seated in the chassis.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Interrupt the boot process **only** after determining the correct timing:

      You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete.

      Only press `Ctrl-C` when you see the message `Press Ctrl-C for Boot Menu`.

      If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort.

      If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode.

      During the boot process, you might see the following prompts:

      - A prompt warning of a system ID mismatch and asking to override the system ID.
      - A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.

You can safely respond `y` to these prompts.

e. Select the option to boot to Maintenance mode from the displayed menu.

Restore and verify the system configuration - AFF C190

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

The value for HA-state can be one of the following:

- ha
- mcc
- mcc-2n
- mccip
- non-ha

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - AFF C190

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Verifying the system ID change on an HA system

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering **`y`** if you are prompted to override the system ID due to a system ID mismatch.
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node                Partner         Possible     State Description
------------      ------------      --------
-------------------------------------
node1               node2           false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2               node1           -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond **`y`** when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter **`y`**.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the `System ID`

`changed on partner` message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. Verify that the expected volumes are present for each controller: `vol show -node node-name`
9. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF C190

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Installing licenses for the *replacement* controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are

invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verifying LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - AFF C190

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove controller module

To access components inside the controller module, you must first remove the controller module from the system, and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

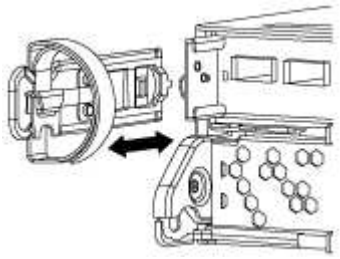

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

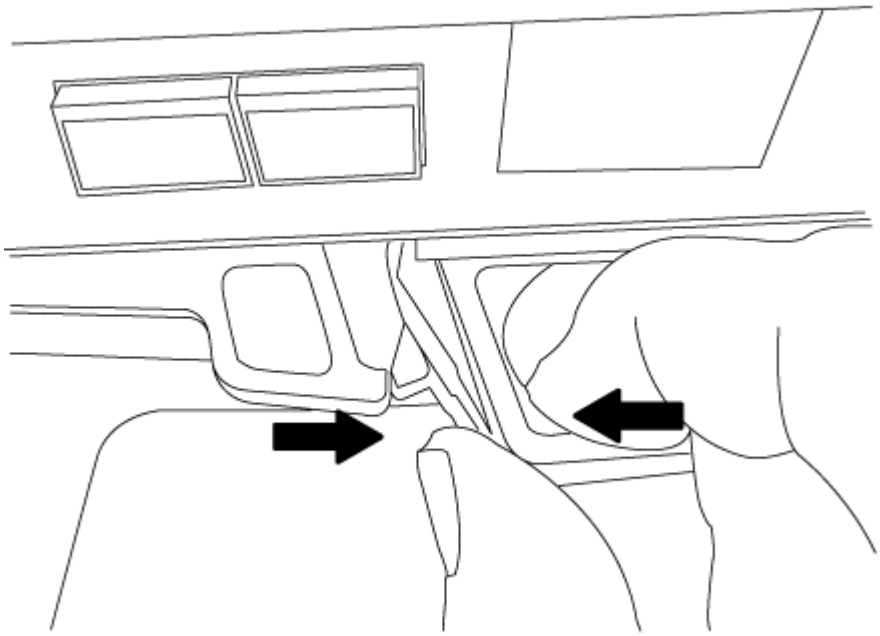

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

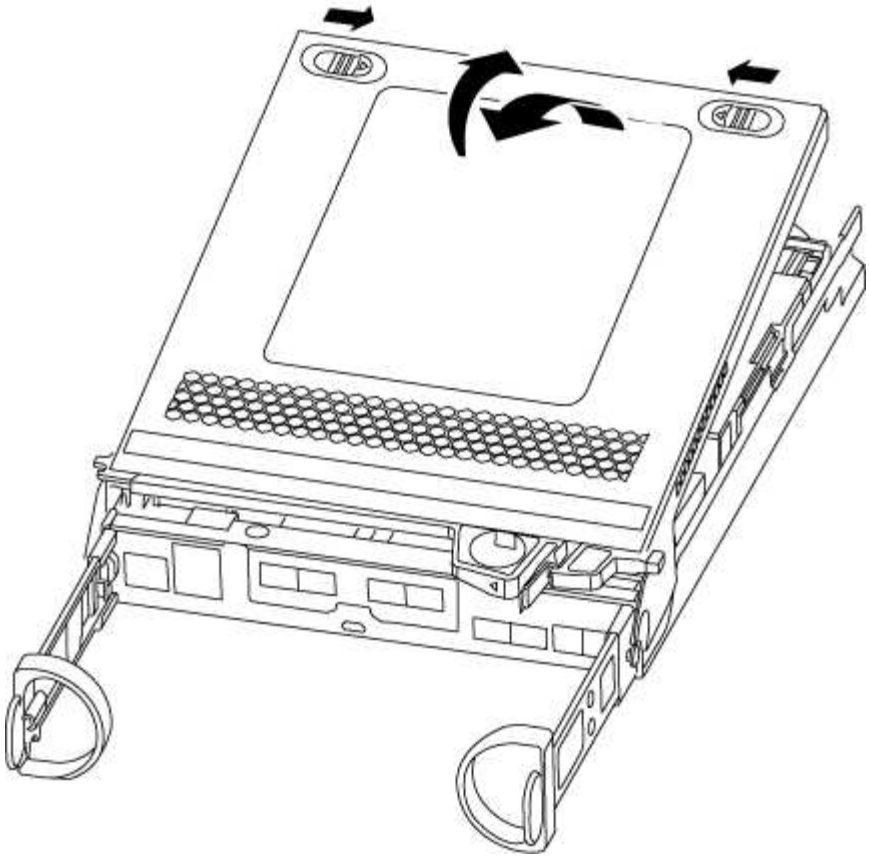

Step 3: Replace the DIMMs

To replace the DIMMs, you need to locate them inside the controller module, and then follow the specific sequence of steps.

If you are replacing a DIMM, you need to remove it after you have unplugged the NVMEM battery from the controller module.

1. Check the NVMEM LED on the controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

   

2. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
3. If the NVMEM LED is flashing, there is data in the NVMEM and you must disconnect the battery to clear the memory:
   a. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

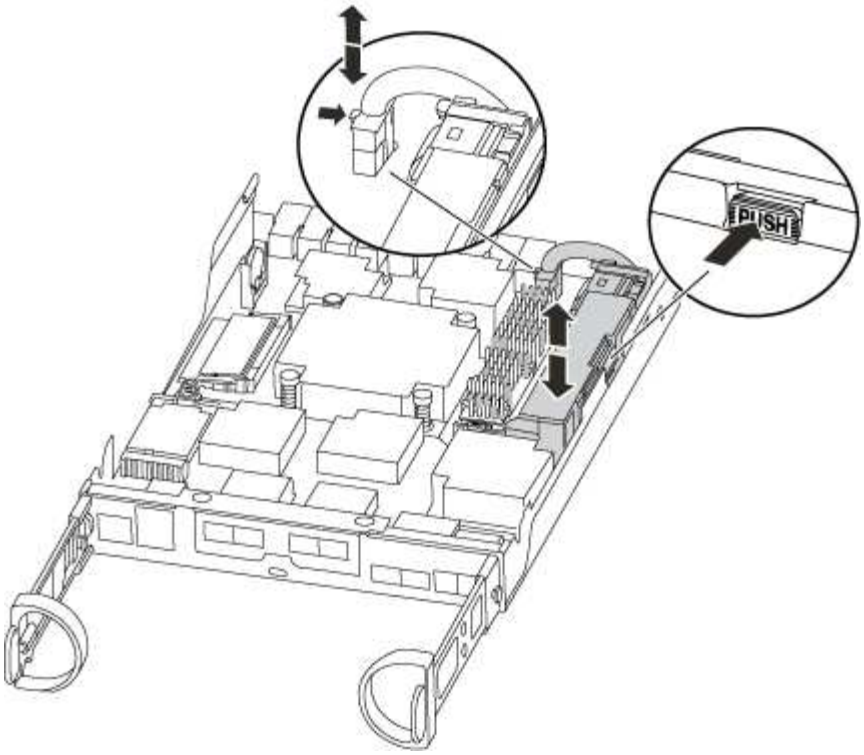

b. Confirm that the NVMEM LED is no longer lit.

c. Reconnect the battery connector.

4. Return to Step 3: Replace the DIMMs in this procedure to recheck the NVMEM LED.
5. Locate the DIMMs on your controller module.
6. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
7. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

> Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

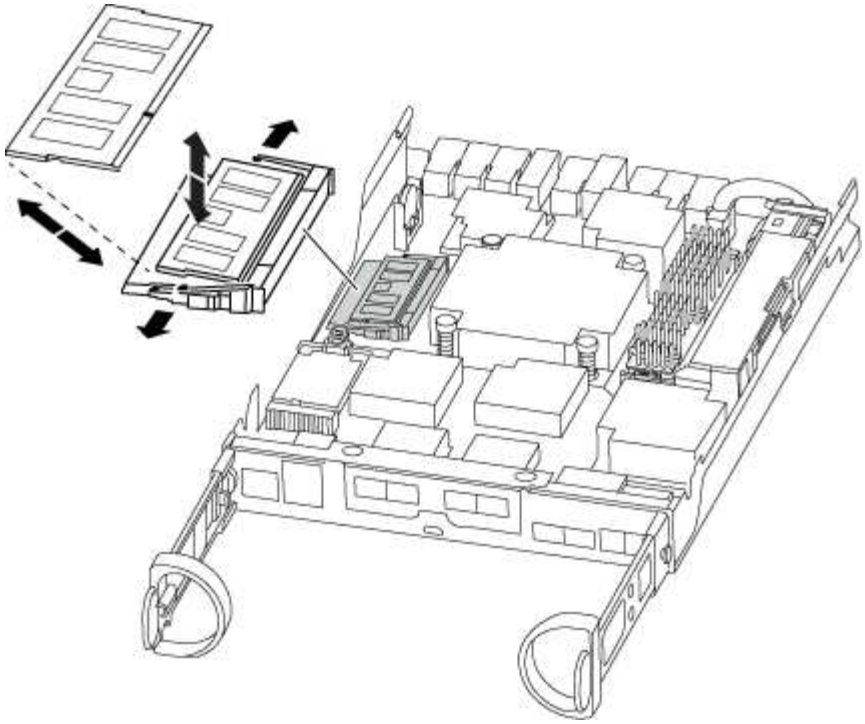

8. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

9. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

10. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
11. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

    Make sure that the plug locks down onto the controller module.

12. Close the controller module cover.

Step 4: Reinstall the controller module

After you replace components in the controller module, you must reinstall it into the chassis.

1. If you have not already done so, replace the cover on the controller module.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module. The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller begins to boot as soon as it is seated in the chassis.

   b. If you have not already done so, reinstall the cable management device.

   c. Bind the cables to the cable management device with the hook and loop strap.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - AFF C190

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace the NVMEM battery - AFF C190

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove controller module

To access components inside the controller module, you must first remove the controller module from the system, and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

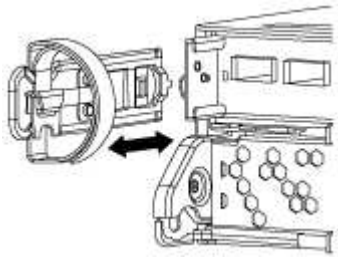

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

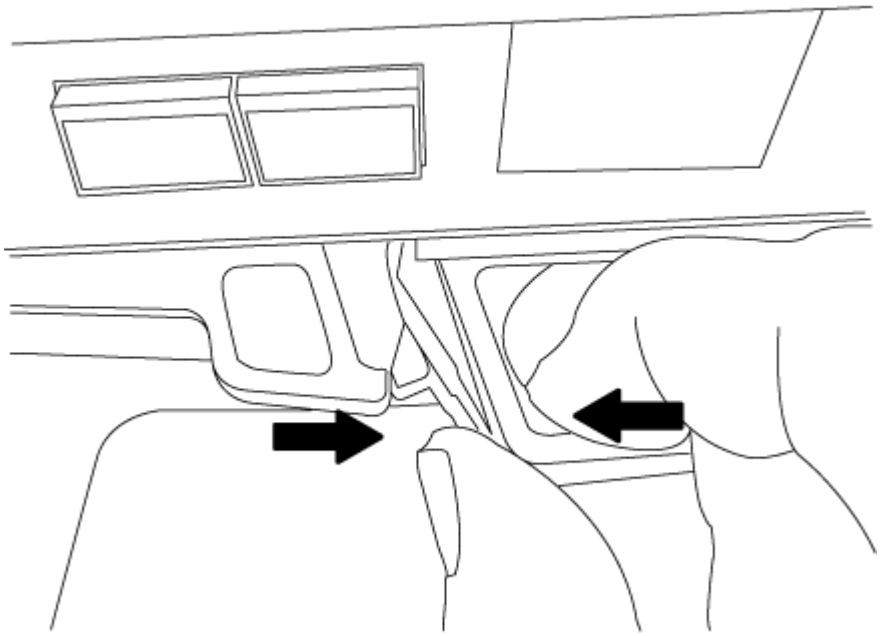

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

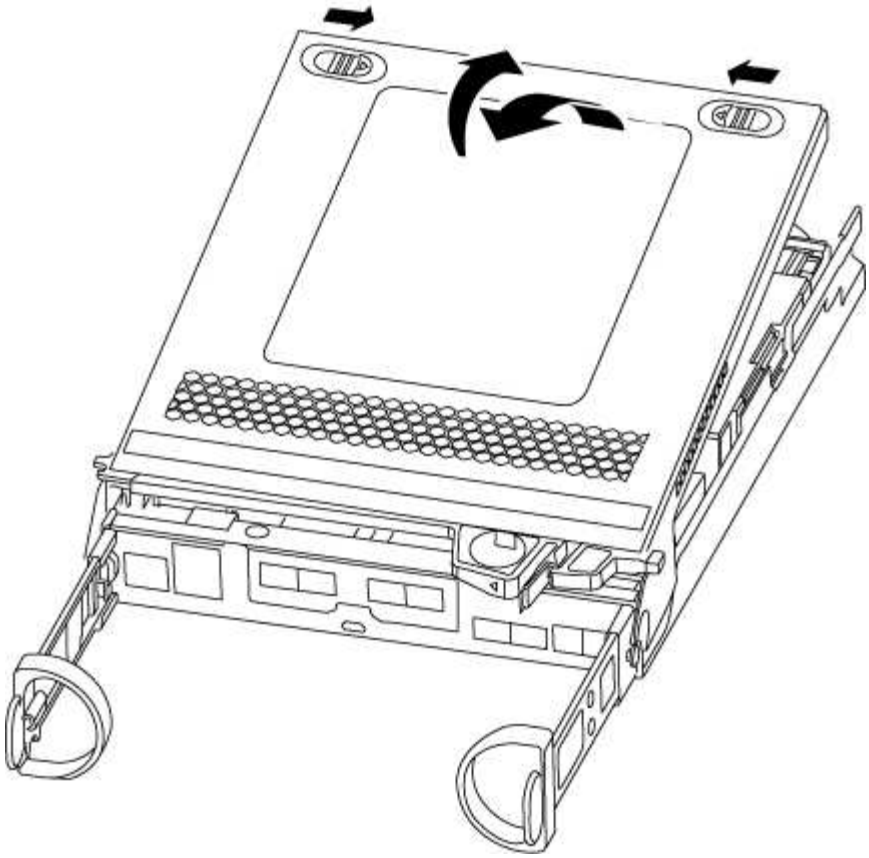

Step 3: Replace the NVMEM battery

To replace the NVMEM battery in your system, you must remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

- If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
- If the LED is on and power is on, unwritten data is stored on NVMEM.

  This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Locate the NVMEM battery in the controller module.

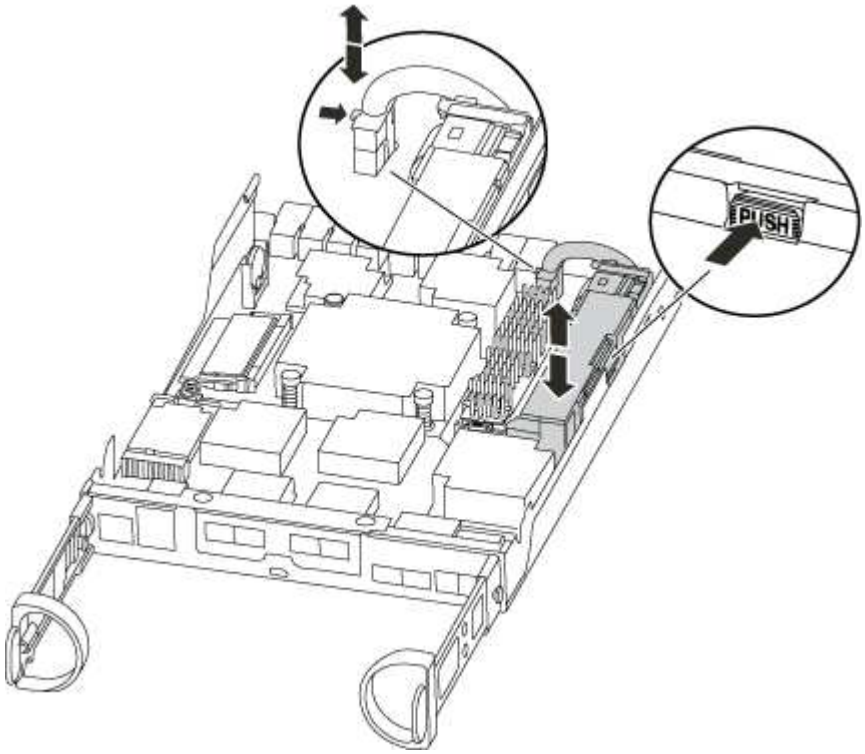

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Remove the battery from the controller module and set it aside.
5. Remove the replacement battery from its package.
6. Loop the battery cable around the cable channel on the side of the battery holder.
7. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
8. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
9. Plug the battery plug back into the controller module.

Step 4: Reinstall the controller module

After you replace components in the controller module, you must reinstall it into the chassis.

1. If you have not already done so, replace the cover on the controller module.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module. The controller module begins to boot as soon as it is fully seated in the chassis.
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller begins to boot as soon as it is seated in the chassis.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - AFF C190

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

Cooling is integrated with the power supply, so you must replace the power supply within two minutes of removal to prevent overheating due to reduced airflow. Because the chassis provides a shared cooling configuration for the two HA nodes, a delay longer than two minutes will shut down all controller modules in the chassis. If both controller modules do shut down, make sure that both power supplies are inserted, turn both off for 30 seconds, and then turn both on.

- Power supplies are auto-ranging.

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.

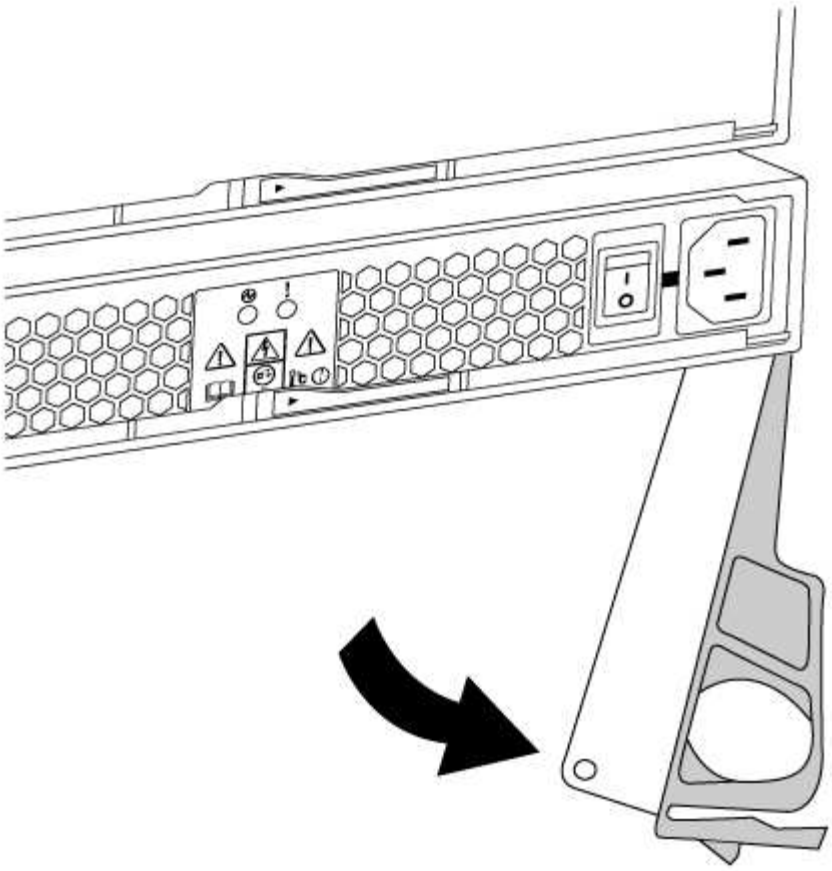

5. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.
7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
9. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

Once power is restored to the power supply, the status LED should be green.

1. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

   The power supply LEDs are lit when the power supply comes online.

2. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Remove controller module

To access components inside the controller module, you must first remove the controller module from the

system, and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

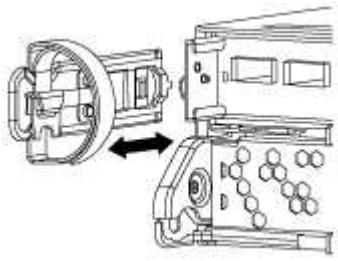

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

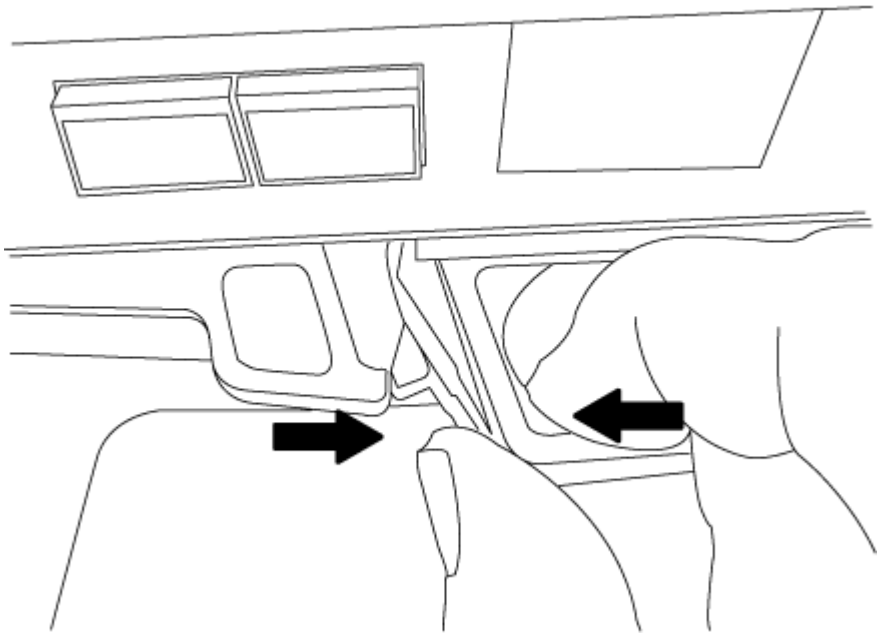

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

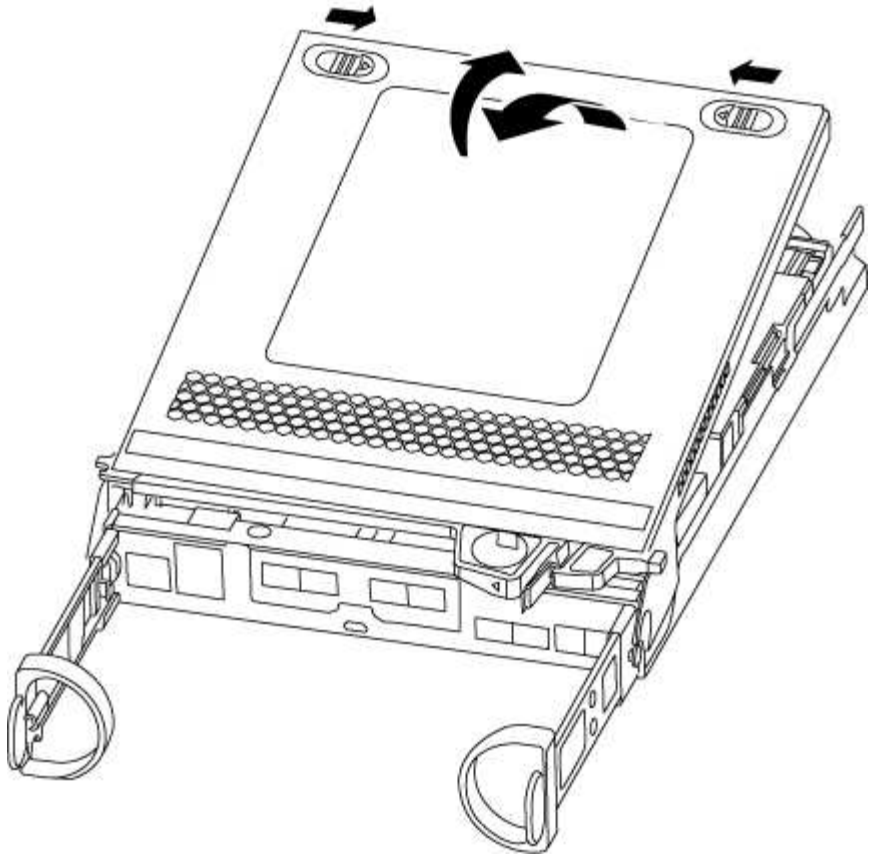

Step 3: Replace the RTC battery

To replace the RTC battery, you need to locate it inside the controller module, and then follow the specific sequence of steps.

1. Locate the RTC battery.

2. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

3. Remove the replacement battery from the antistatic shipping bag.
4. Locate the empty battery holder in the controller module.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and set time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module. The controller module begins to boot as soon as it is fully seated in the chassis.
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Complete the replacement process

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

FAS2600 systems

Install and setup

Cluster configuration worksheet - FAS2600

You can use the worksheet to gather and record your site-specific IP addresses and other information required when configuring an ONTAP cluster.

Cluster Configuration Worksheet

Start here: Choose your installation and setup experience

You can choose from different content formats to guide you through installing and setting up your new storage system.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

Installation and setup PDF poster - FAS2600

You can use the PDF poster to install and set up your new system. The FAS2600 Installation and Setup Instructions provides step-by-step instructions with live links to additional content.

Installation and setup video - FAS2600

The following video shows end-to-end software configuration for systems running ONTAP 9.2.

AFF FAS2600 Setup Video

Maintain

Maintain FAS2600 hardware

For the FAS2600 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Caching module

You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

NVEM battery

A battery is included with a controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - FAS2600

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.

- It is important that you apply the commands in these steps on the correct node:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy node* is the HA partner of the impaired node.

Check encryption key support and status - FAS2600

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.

ONTAP version	Run this command
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the impaired controller - FAS2600

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller displays…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Replace the boot media - FAS2600

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

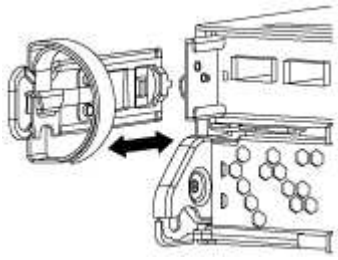

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

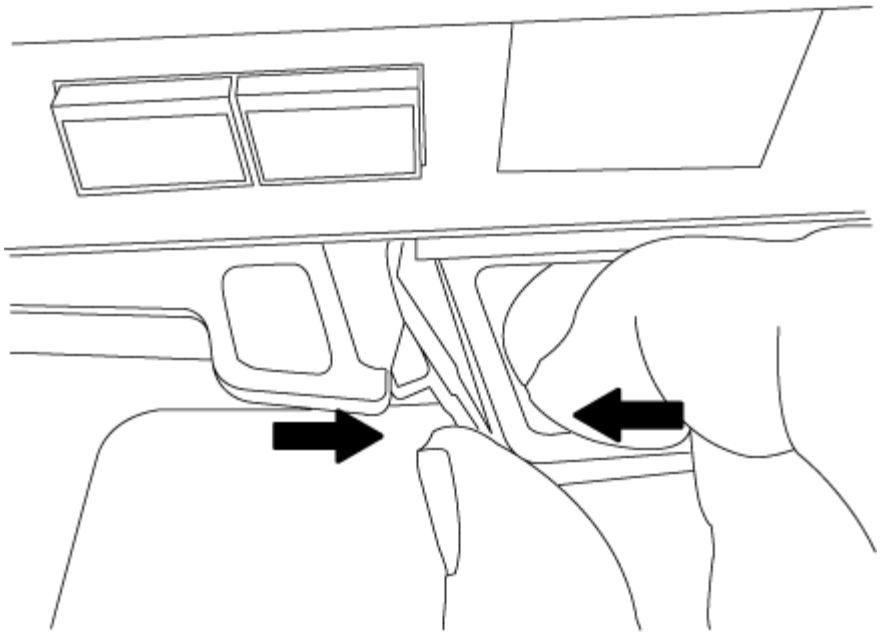

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

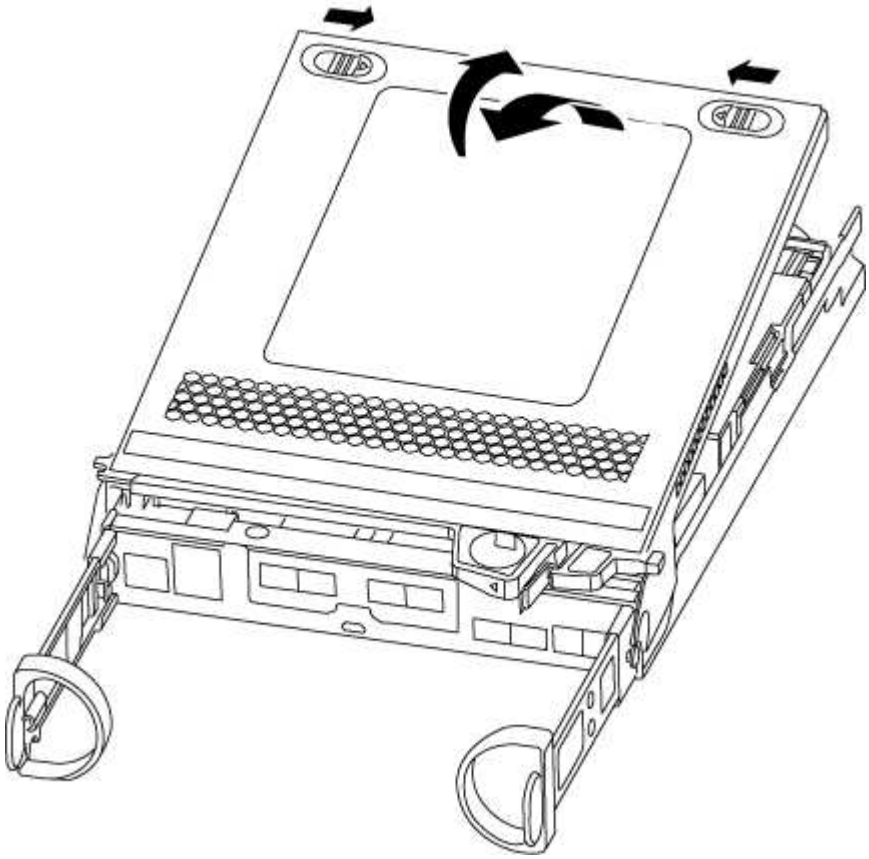

Step 2: Replace the boot media

1. If you are not already grounded, properly ground yourself.
2. Locate the boot media using the following illustration or the FRU map on the controller module:

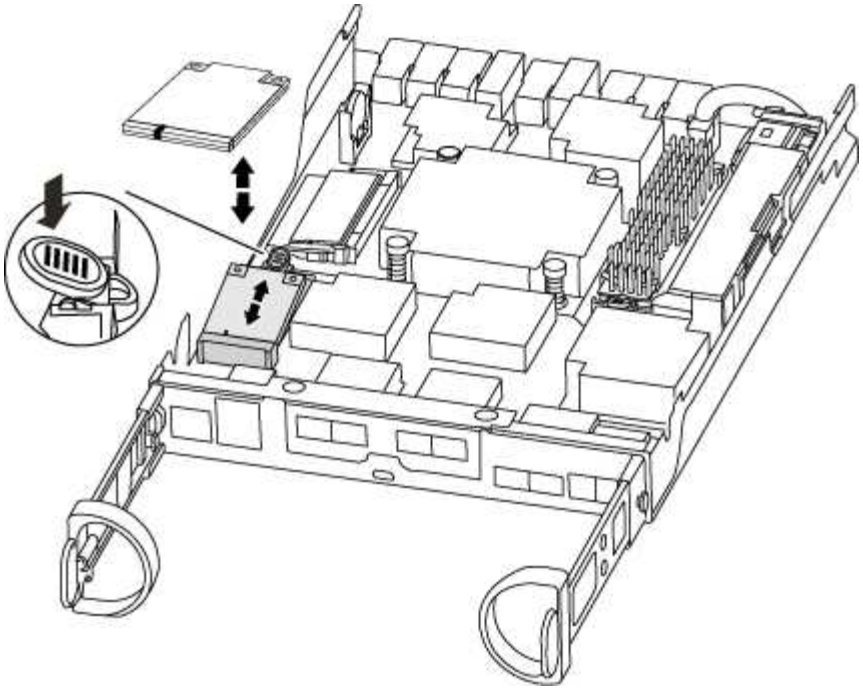

3. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

4. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
5. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

6. Push the boot media down to engage the locking button on the boot media housing.
7. Close the controller module cover.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

   The controller begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the controller to boot to LOADER.

6. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

   The system begins to boot and stops at the LOADER prompt.

7. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

     The target port you configure is the target port you use to communicate with the impaired controller from the healthy controller during var file system restore with a network connection. You can also use the e0M port in this command.

   - If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
     - filer_addr is the IP address of the storage system.
     - netmask is the network mask of the management network that is connected to the HA partner.
     - gateway is the gateway for the network.
     - dns_addr is the IP address of a name server on your network.
     - dns_domain is the Domain Name System (DNS) domain name.

       If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

     Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

Boot the recovery image - FAS2600

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment_variable_name changed_value`* command.
   d. Save your changes using the `saveenv` command.
5. The next depends on your system configuration:
   - If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed
   - If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.
8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore encryption - FAS2600

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - FAS2600

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the caching module - FAS2600

You must replace the caching module in the controller module when your system registers a single AutoSupport (ASUP) message that the module has gone offline; failure to do so results in performance degradation.

- You must replace the failed component with a replacement FRU component you received from your provider.

AFF FAS2600 caching module replacement video

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

About this task

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller.

Synchronize a node with the cluster

You might want to erase the contents of your caching module before replacing it.

1. Although data on the caching module is encrypted, you might want to erase any data from the impaired caching module and verify that the caching module has no data:
   a. Erase the data on the caching module: `system controller flash-cache secure-erase run -node node name localhost -device-id device_number`

      Run the `system controller flash-cache show` command if you don't know the Flash Cache device ID.

   b. Verify that the data has been erased from the caching module: `system controller flash-cache secure-erase show`
2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller: • For an HA pair, take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`. • For a stand-alone system: `system node halt impaired_node_name`

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

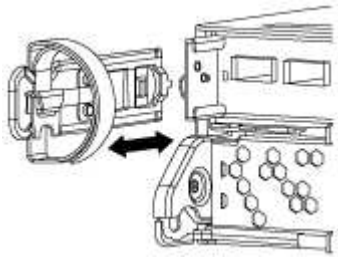

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

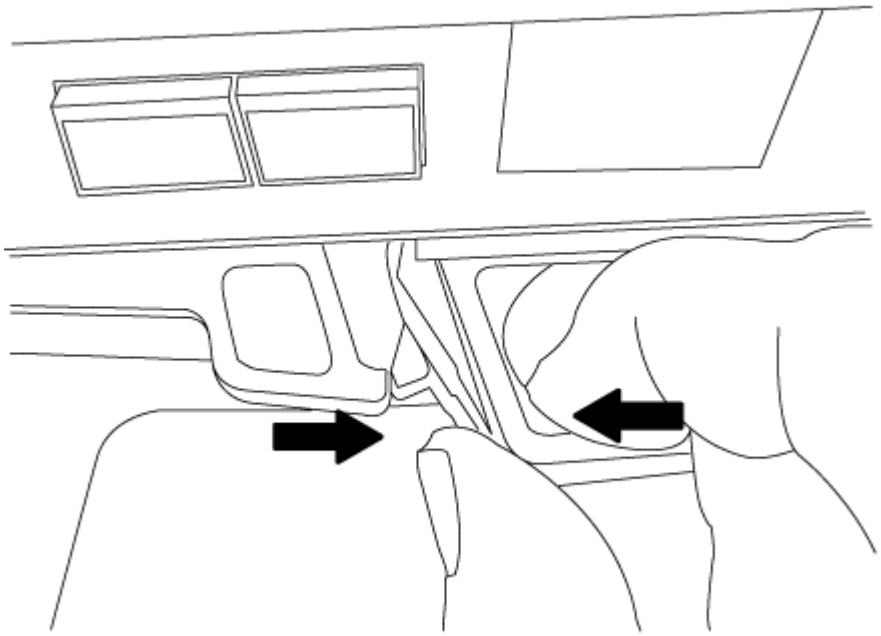

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

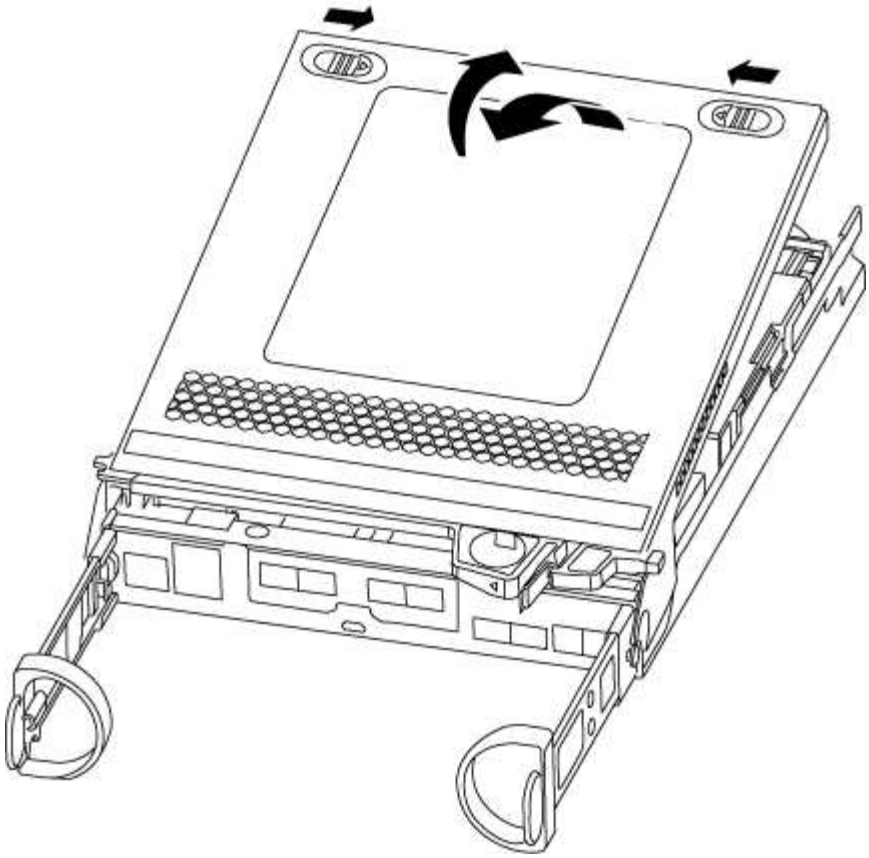

Step 3: Replace a caching module

To replace a caching module referred to as the M.2 PCIe card on the label on your controller, locate the slot inside the controller and follow the specific sequence of steps.

Your storage system must meet certain criteria depending on your situation:

- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- All other components in the storage system must be functioning properly; if not, you must contact technical support.

1. Locate the caching module at the rear of the controller module and remove it.
   a. Press the release tab.
   b. Remove the heatsink.

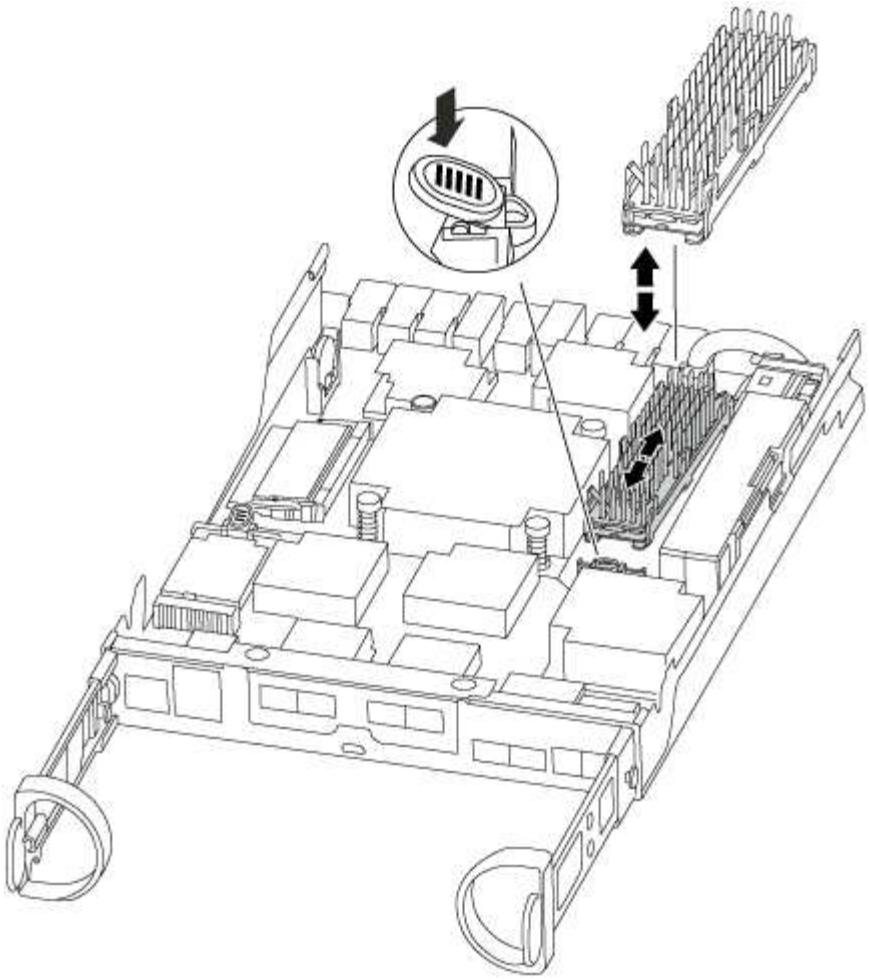

1. Gently pull the caching module straight out of the housing.
2. Align the edges of the caching module with the socket in the housing, and then gently push it into the socket.
3. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

4. Reseat and push the heatsink down to engage the locking button on the caching module housing.
5. Close the controller module cover, as needed.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

1. If you have not already done so, replace the cover on the controller module.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber

optic cables.

4. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, then turn on the power to start the boot process.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - FAS2600

To replace the chassis, you must move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving all drives and controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

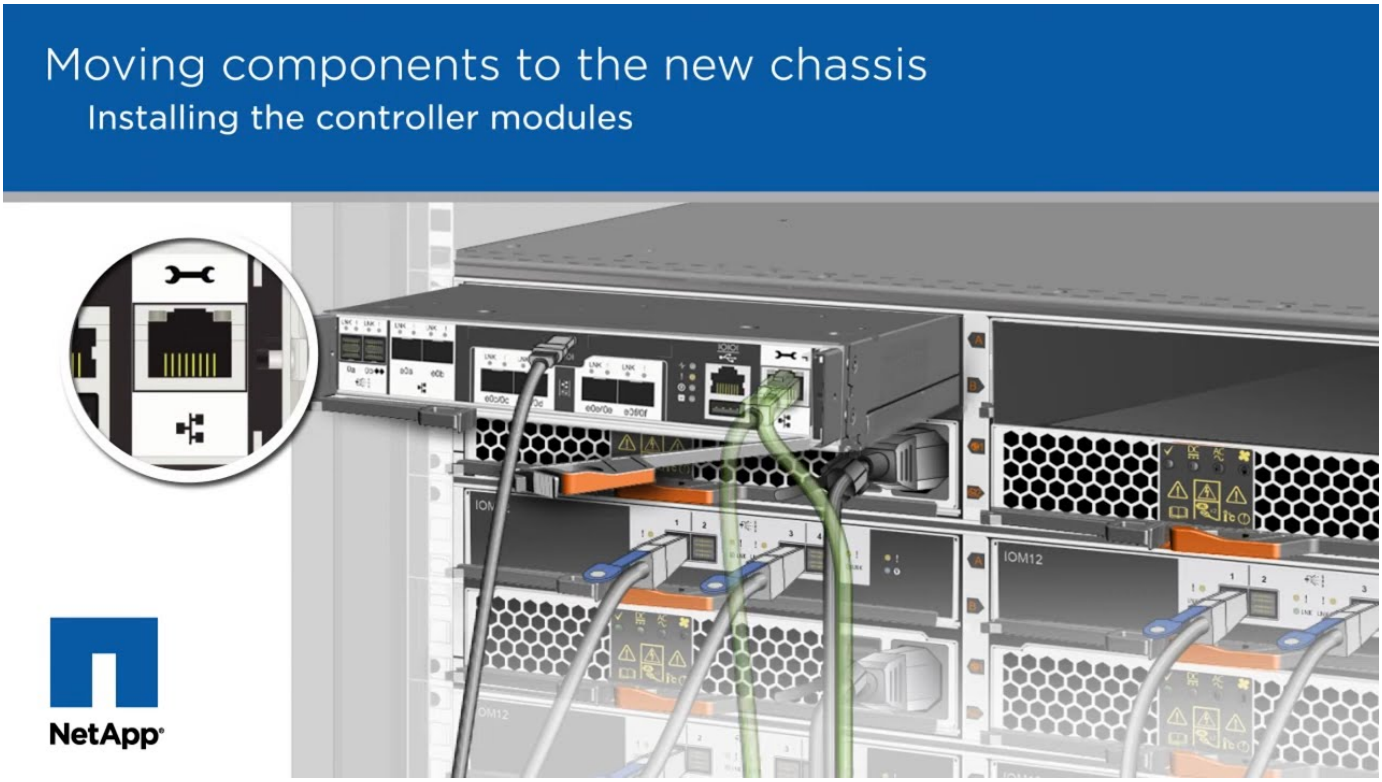

Shut down the controllers - FAS2600

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.

- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - FAS2600

Move the power supplies, hard drives, and controller module or modules from the

impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Move the power supply

Moving out a power supply when replacing a chassis involves turning off, disconnecting, and removing the power supply from the old chassis and installing and connecting it on the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.
4. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

   Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

Step 2: Remove the controller module

Remove the controller module or modules from the old chassis.

1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

2. Remove and set aside the cable management devices from the left and right sides of the controller module.

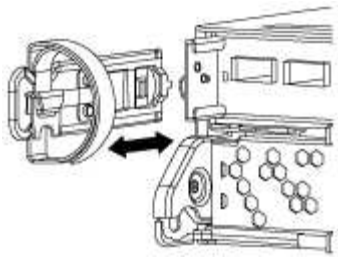

3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

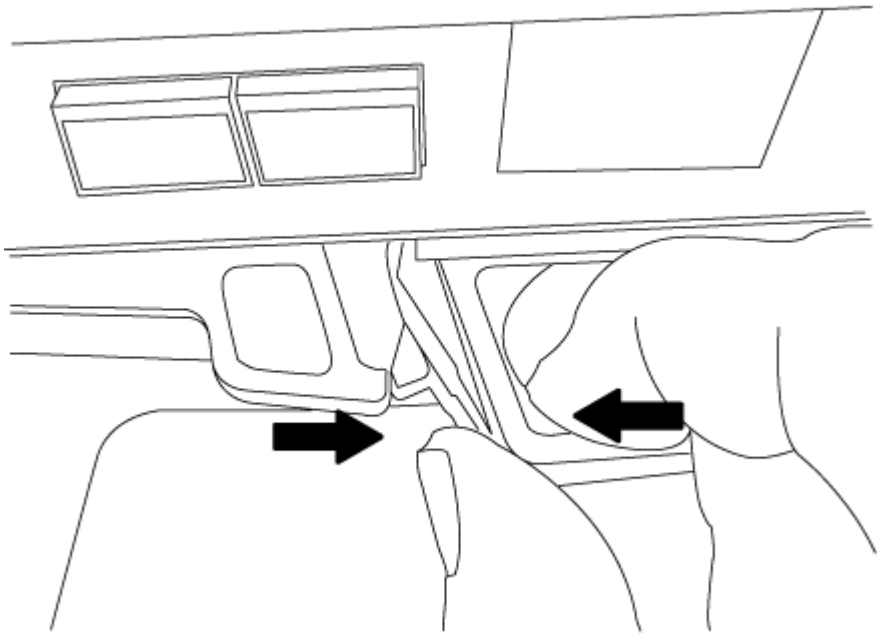

4. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 3: Move drives to the new chassis

Move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.

3. Align the drive from the old chassis with the same bay opening in the new chassis.

4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.

5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 4: Replace a chassis from within the equipment rack or system cabinet

Remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller

After you install the controller module and any other components into the new chassis, boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Repeat the preceding steps if there is a second controller to install in the new chassis.
4. Complete the installation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Repeat the preceding steps for the second controller module in the new chassis.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reinstall the blanking panel and then go to the next step.

5. Connect the power supplies to different power sources, and then turn them on.
6. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

      If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

   b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - FAS2600

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis *HA-state*`

      The value for HA-state can be one of the following:

      - `ha`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. The next step depends on your system configuration.

If your system is in…	Then…
A stand-alone configuration	a. Exit Maintenance mode: `halt` b. Go to Completing the replacement process.
An HA pair with a second controller module	Exit Maintenance mode: `halt` The LOADER prompt appears.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller module

Overview of controller module replacement - FAS2600

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- This procedure includes steps for automatically or manually reassigning drives to the *replacement* controller, depending on your system’s configuration.

You should perform the drive reassignment as directed in the procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the controller - FAS2600

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Replace the controller module hardware - FAS2600

To replace the controller module, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

AFF FAS2600 controller replacement video

Step 1: Remove controller module

To replace the controller module, you must first remove the old controller module from the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

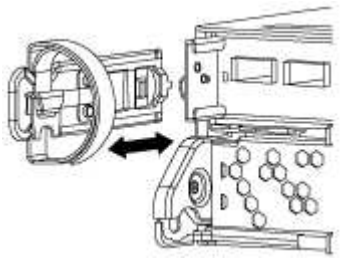

4. If you left the SFP modules in the system after removing the cables, move them to the new controller module.
5. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

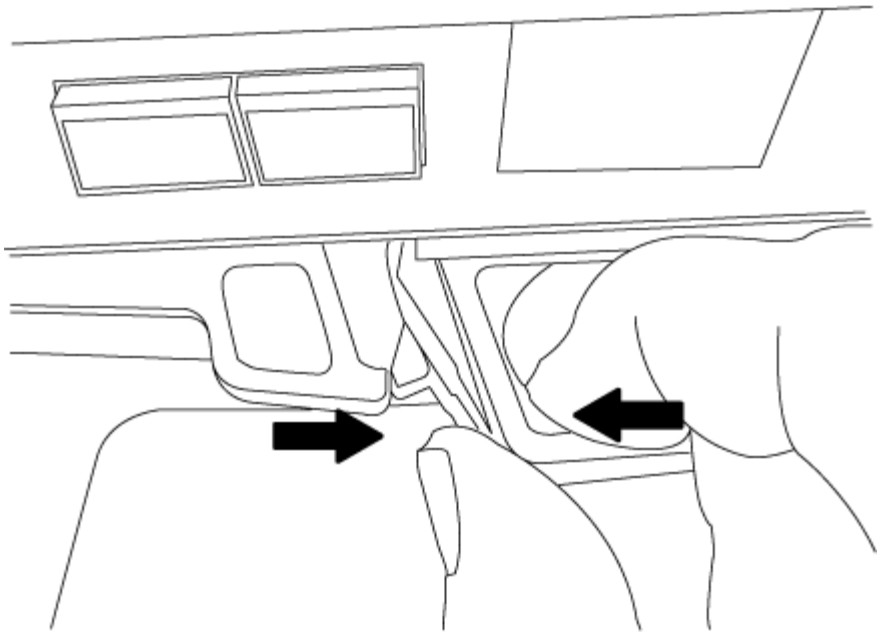

6. Turn the controller module over and place it on a flat, stable surface.
7. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

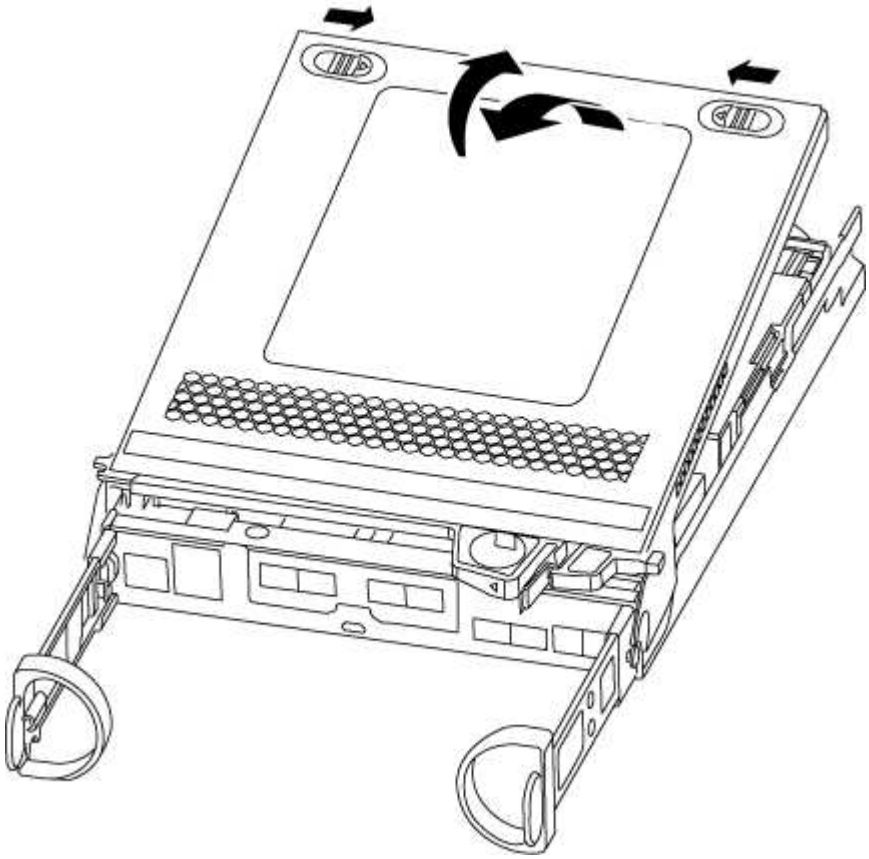

Step 2: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller module and insert it in the new controller module.

Steps

1. Locate the boot media using the following illustration or the FRU map on the controller module:

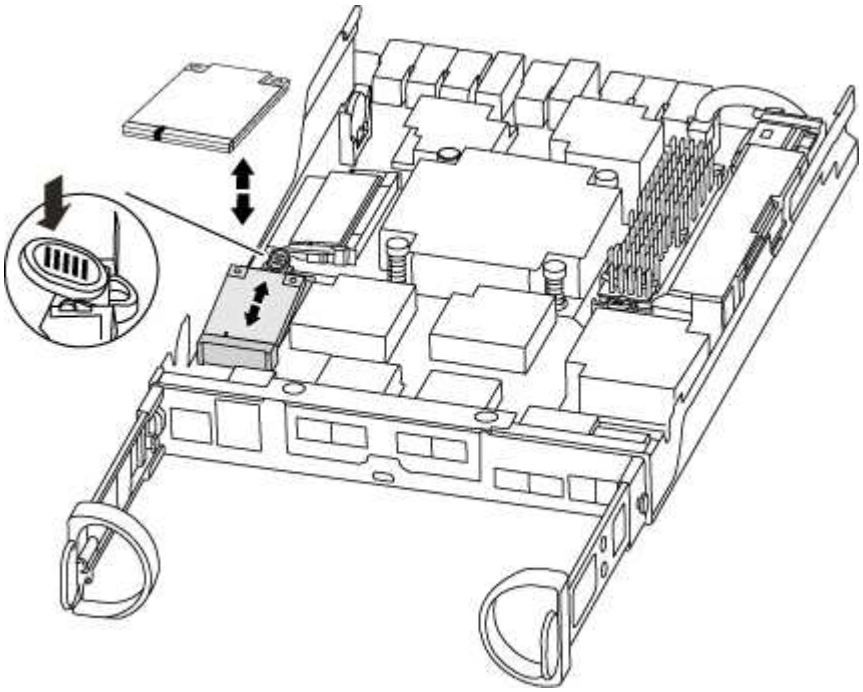

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the NVMEM battery

To move the NVMEM battery from the old controller module to the new controller module, you must perform a specific sequence of steps.

Steps

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

- If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
- If the LED is on and power is on, unwritten data is stored on NVMEM.

  This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Locate the NVMEM battery in the controller module.

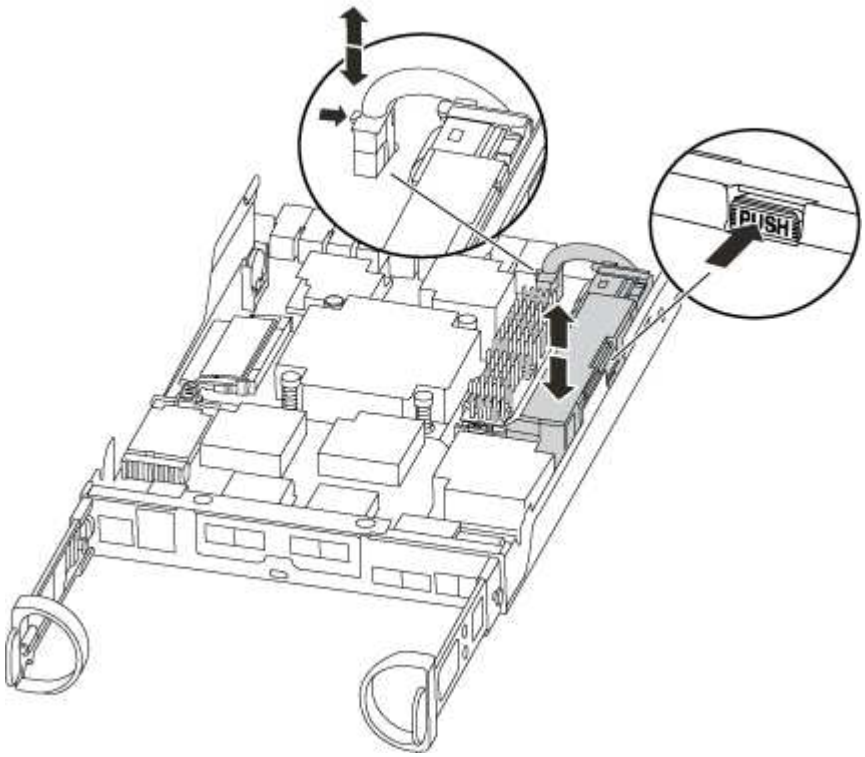

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Loop the battery cable around the cable channel on the side of the battery holder.
7. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
8. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.

Step 4: Move the DIMMs

To move the DIMMs, you must follow the directions to locate and move them from the old controller module into the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

Steps

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

   The number and placement of system DIMMs depends on the model of your system.

   The following illustration shows the location of system DIMMs:

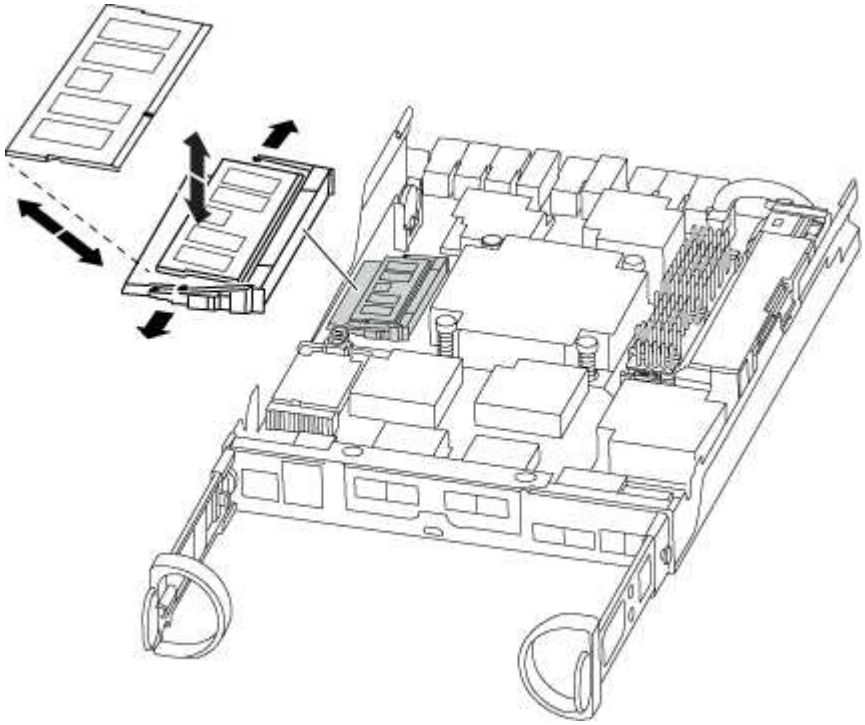

4. Repeat these steps to remove additional DIMMs as needed.
5. Verify that the NVMEM battery is not plugged into the new controller module.
6. Locate the slot where you are installing the DIMM.
7. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Repeat these steps for the remaining DIMMs.
9. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

   Make sure that the plug locks down onto the controller module.

Step 5: Move the caching module

To move a caching module referred to as the M.2 PCIe card on the label on your controller, locate and move it from the old controller into the replacement controller and follow the specific sequence of steps.

You must have the new controller module ready so that you can move the caching module directly from the old controller module to the corresponding slot in the new one. All other components in the storage system must be functioning properly; if not, you must contact technical support.

Steps

1. Locate the caching module at the rear of the controller module and remove it.

   a. Press the release tab.

   b. Remove the heatsink.

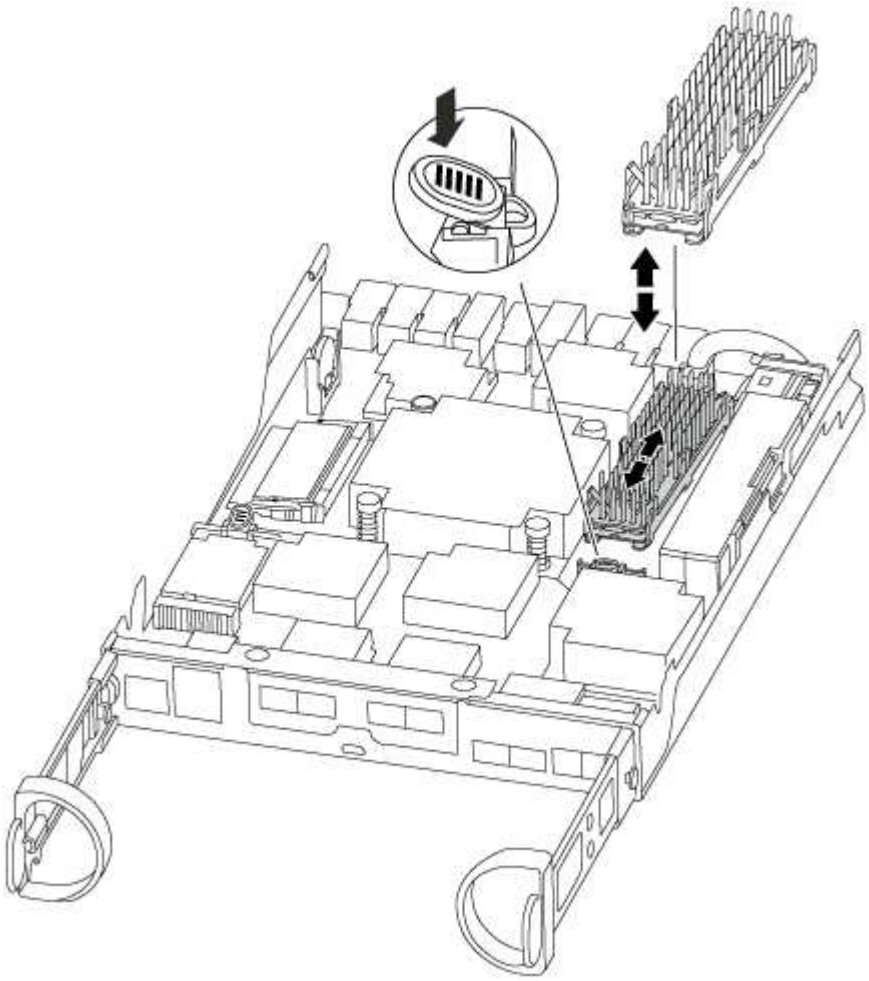

2. Gently pull the caching module straight out of the housing.
3. Move the caching module to the new controller module, and then align the edges of the caching module with the socket housing and gently push it into the socket.
4. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

5. Reseat and push the heatsink down to engage the locking button on the caching module housing.
6. Close the controller module cover, as needed.

Step 6: Install the controller

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

> The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   > Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   > You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. When you see the message `Press Ctrl-C for Boot Menu`, press `Ctrl-C` to interrupt the boot process. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. Select the option to boot to Maintenance mode from the displayed menu.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process, and then press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. From the boot menu, select the option for Maintenance mode.

Important: During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
  You can safely respond `y` to these prompts.

Restore and verify the system configuration - FAS2600

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - FAS2600

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks.
You must use the correct procedure for your configuration.

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch: `boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:

a. Change to the advanced privilege level: `set -privilege advanced`

   You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

b. Save any coredumps: `system node run -node *local-node-name* partner savecore`

c. Wait for the `savecore`command to complete before issuing the giveback.

   You can enter the following command to monitor the progress of the savecore command: `system node run -node *local-node-name* partner savecore -s`

d. Return to the admin privilege level: `set -privilege admin`

5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode *replacement_node_name*`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. Verify that the expected volumes are present for each controller: `vol show -node node-name`
9. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Option 2: Manually reassign the system ID on a stand-alone system in ONTAP

In a stand-alone system, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems that are in a stand-alone configuration.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by pressing Ctrl-C, and then select the option to boot to Maintenance mode from the displayed menu.
2. You must enter `Y` when prompted to override the system ID due to a system ID mismatch.
3. View the system IDs: `disk show -a`
4. You should make a note of the old system ID, which is displayed as part of the disk owner column.

   The following example shows the old system ID of 118073209:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                    POOL   SERIAL NUMBER  HOME
--------    -------------            -----  -------------  -------------
disk_name    system-1  (118073209)   Pool0  J8XJE9LC       system-1
(118073209)
disk_name    system-1  (118073209)   Pool0  J8Y478RC       system-1
(118073209)
.
.
.
```

5. Reassign disk ownership by using the system ID information obtained from the disk show command: `disk reassign -s old system ID disk reassign -s 118073209`
6. Verify that the disks were assigned correctly: `disk show -a`

   The disks belonging to the replacement node should show the new system ID. The following example now show the disks owned by system-1 the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                    POOL   SERIAL NUMBER  HOME
--------    -------------            -----  -------------  -------------
disk_name    system-1  (118065481)   Pool0  J8Y0TDZC       system-1
(118065481)
disk_name    system-1  (118065481)   Pool0  J8Y0TDZC       system-1
(118065481)
.
.
.
```

7. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
8. Boot the node: `boot_ontap`

Complete system restoration - FAS2600

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the

failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver`

`* -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS2600

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

AFF FAS2600 DIMM replacement video

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy

controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

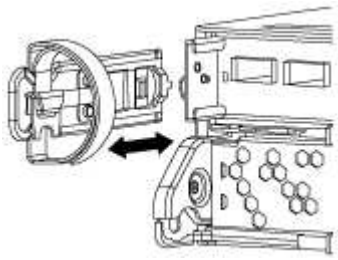

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

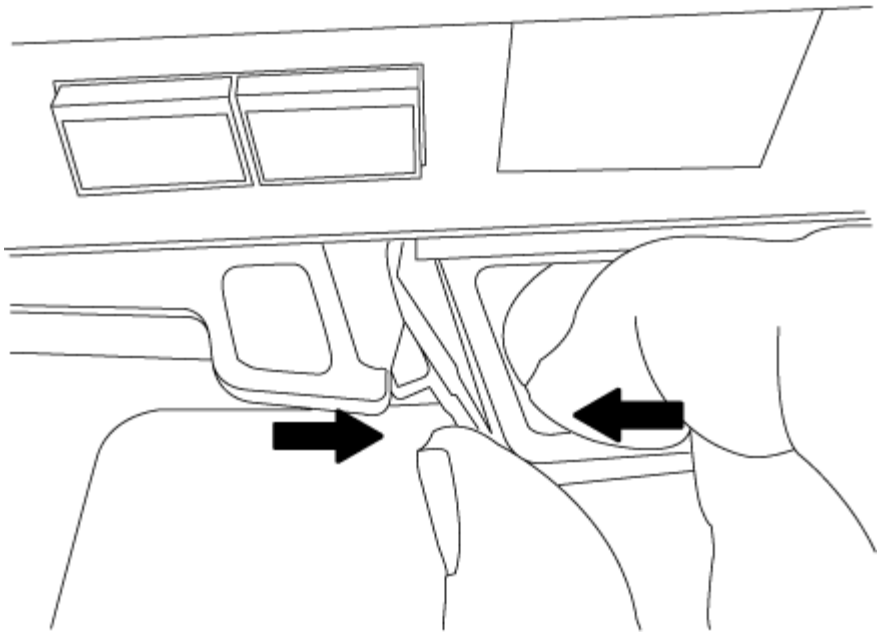

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

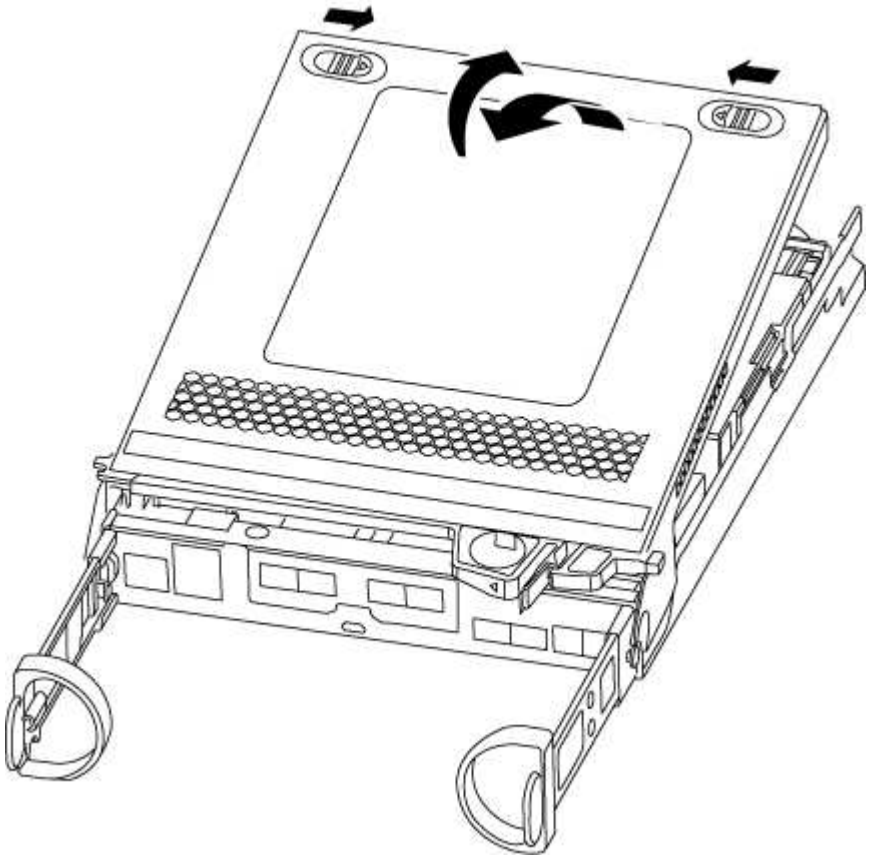

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

If you are replacing a DIMM, you need to remove it after you have unplugged the NVMEM battery from the controller module.

1. Check the NVMEM LED on the controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

   

2. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
3. If the NVMEM LED is flashing, there is data in the NVMEM and you must disconnect the battery to clear the memory:
   a. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

   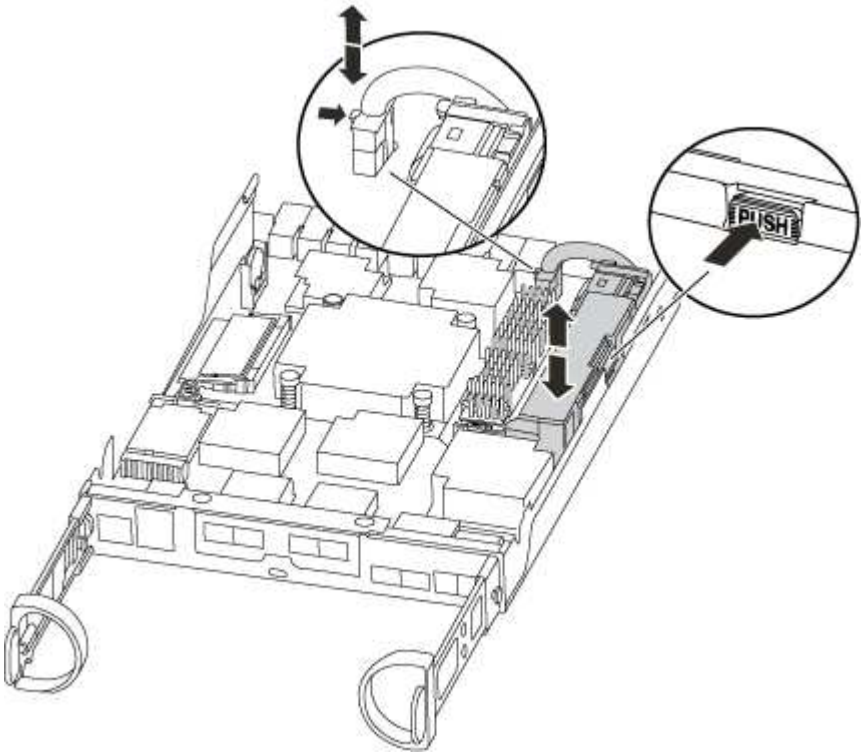

   b. Confirm that the NVMEM LED is no longer lit.
   c. Reconnect the battery connector.
4. Recheck the NVMEM LED.
5. Locate the DIMMs on your controller module.
6. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
7. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

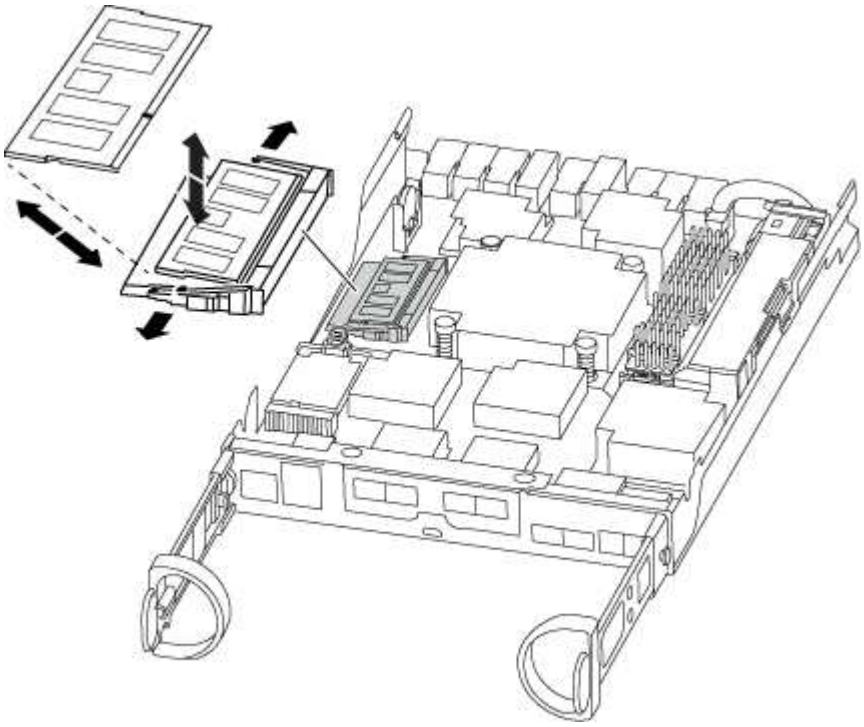

8. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

9. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

10. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
11. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

    Make sure that the plug locks down onto the controller module.

12. Close the controller module cover.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

1. If you have not already done so, replace the cover on the controller module.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, then turn on the power to start the boot process.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - FAS2600

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  > Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace the NVMEM battery - FAS2600

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

AFF FAS2600 NVMEM battery replacement video

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

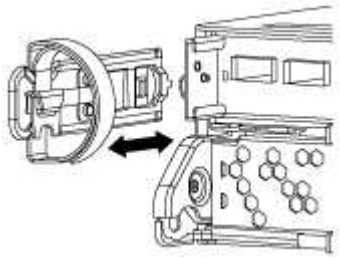

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

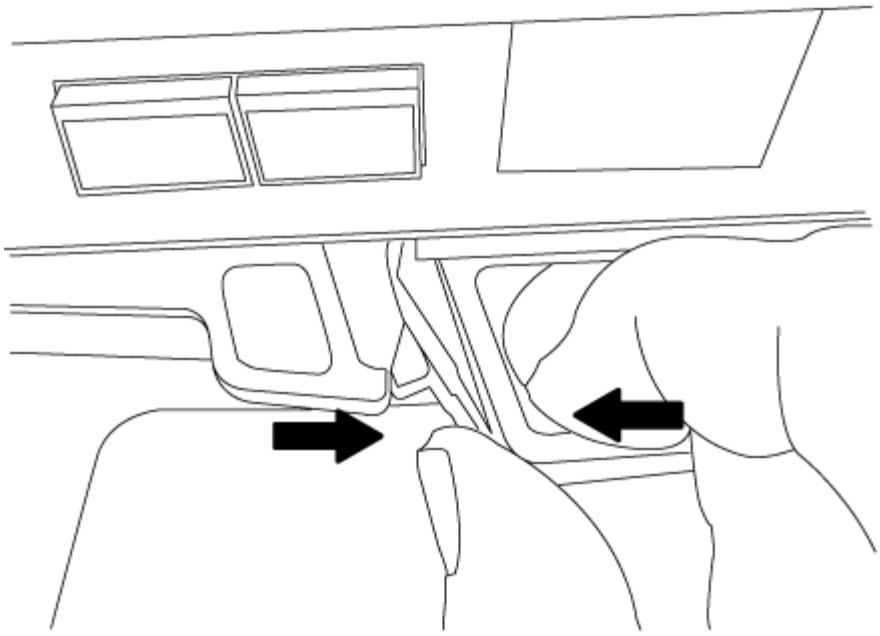

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

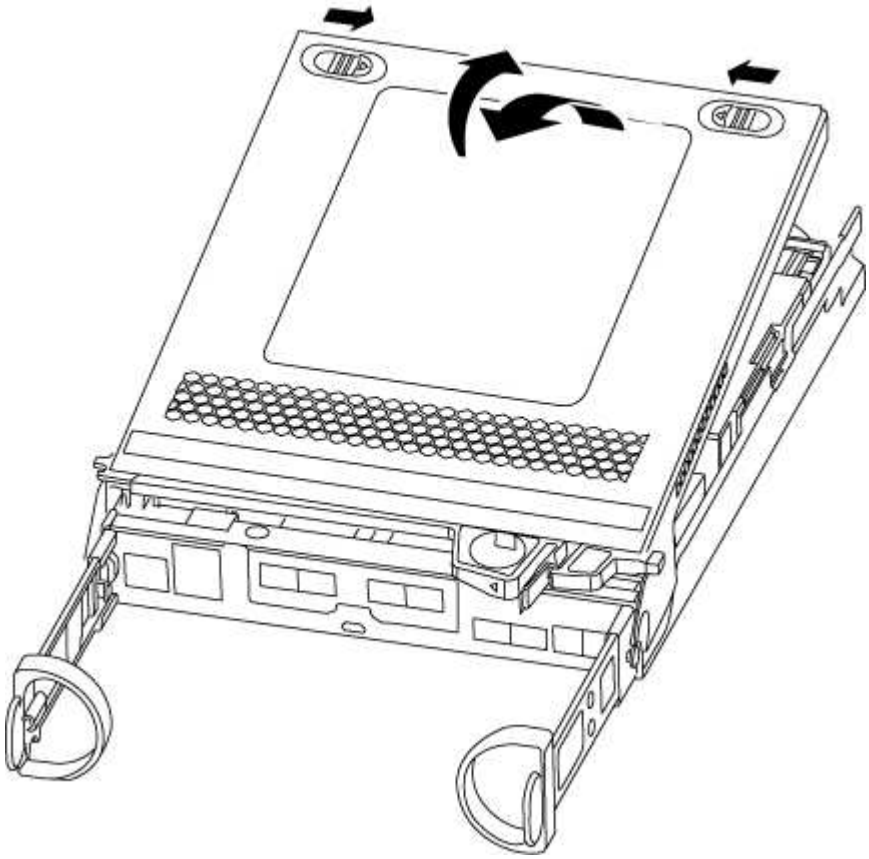

Step 3: Replace the NVMEM battery

To replace the NVMEM battery in your system, you must remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Locate the NVMEM battery in the controller module.

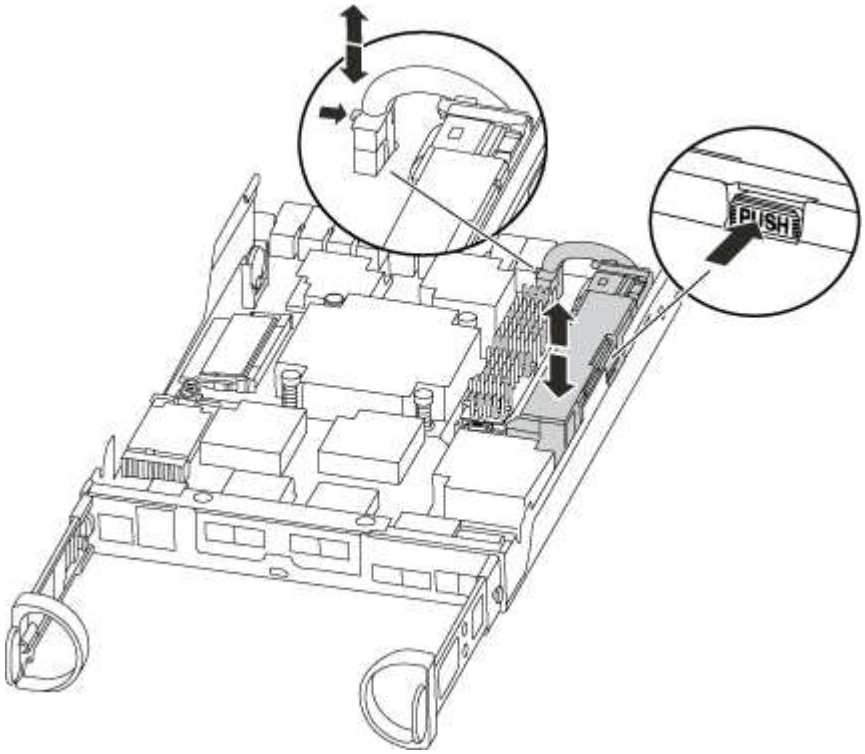

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Remove the battery from the controller module and set it aside.
5. Remove the replacement battery from its package.
6. Loop the battery cable around the cable channel on the side of the battery holder.
7. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.

8. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
9. Plug the battery plug back into the controller module.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

1. If you have not already done so, replace the cover on the controller module.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and turn on the power to start the boot process.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - FAS2600

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> Cooling is integrated with the power supply, so you must replace the power supply within two minutes of removal to prevent overheating due to reduced airflow. Because the chassis provides a shared cooling configuration for the two HA nodes, a delay longer than two minutes will shut down all controller modules in the chassis. If both controller modules do shut down, make sure that both power supplies are inserted, turn both off for 30 seconds, and then turn both on.

- The number of power supplies in the system depends on the model.
- Power supplies are auto-ranging.

AFF FAS2600 power supply replacement video

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:

a. Turn off the power switch on the power supply.

b. Open the power cable retainer, and then unplug the power cable from the power supply.

c. Unplug the power cable from the power source.

4. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.

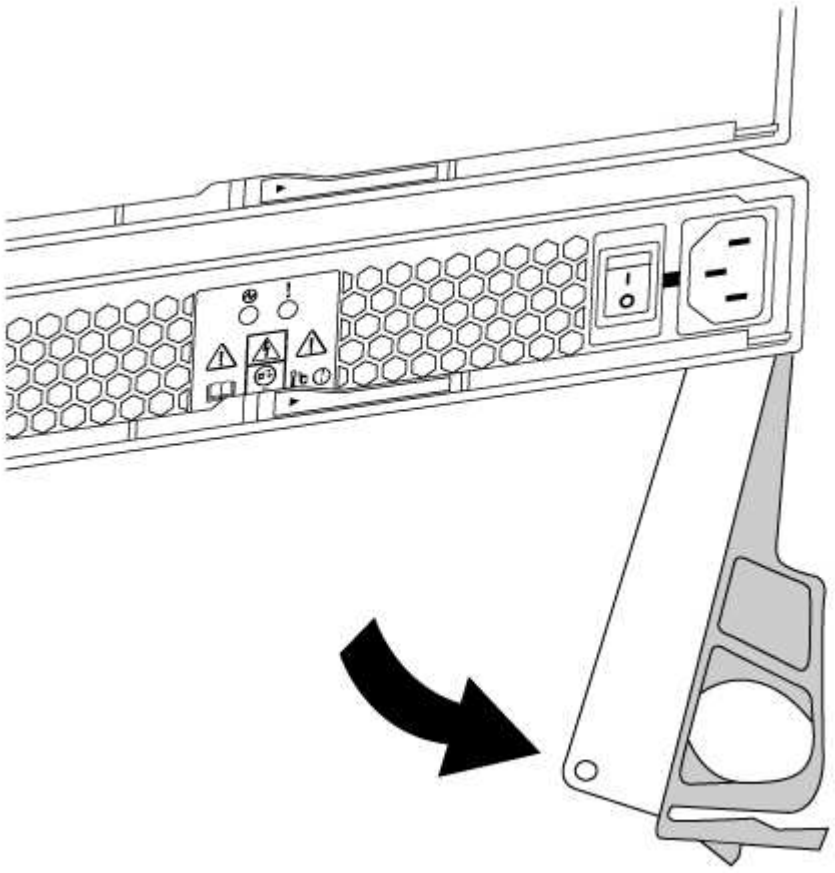

5. Use the cam handle to slide the power supply out of the system.

When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.

7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

The power supplies are keyed and can only be installed one way.

Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.

9. Reconnect the power supply cabling:

a. Reconnect the power cable to the power supply and the power source.

b. Secure the power cable to the power supply using the power cable retainer.

Once power is restored to the power supply, the status LED should be green.

10. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

    The power supply LEDs are lit when the power supply comes online.

11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

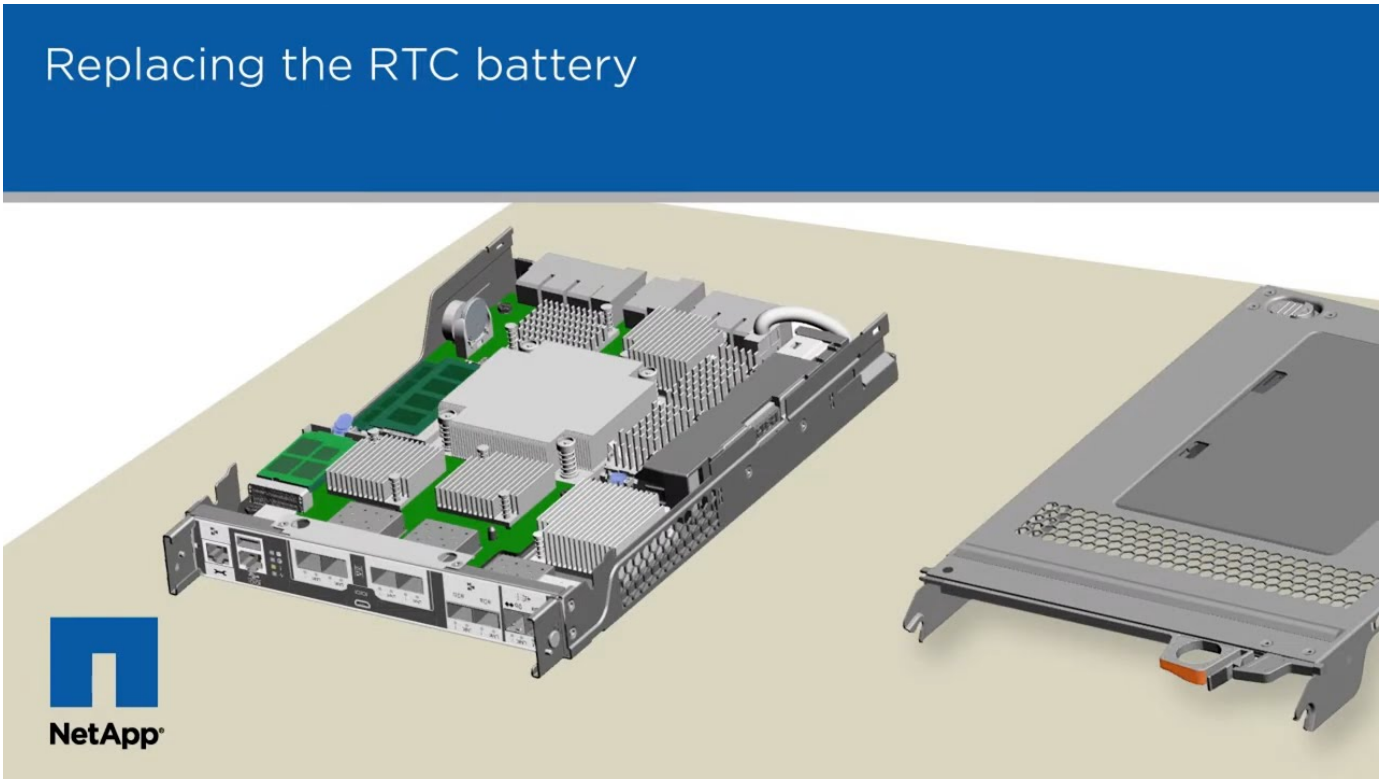

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	**Then…**
The LOADER prompt	Go to Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

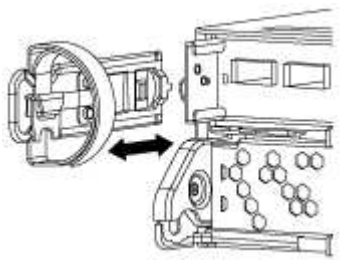

4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

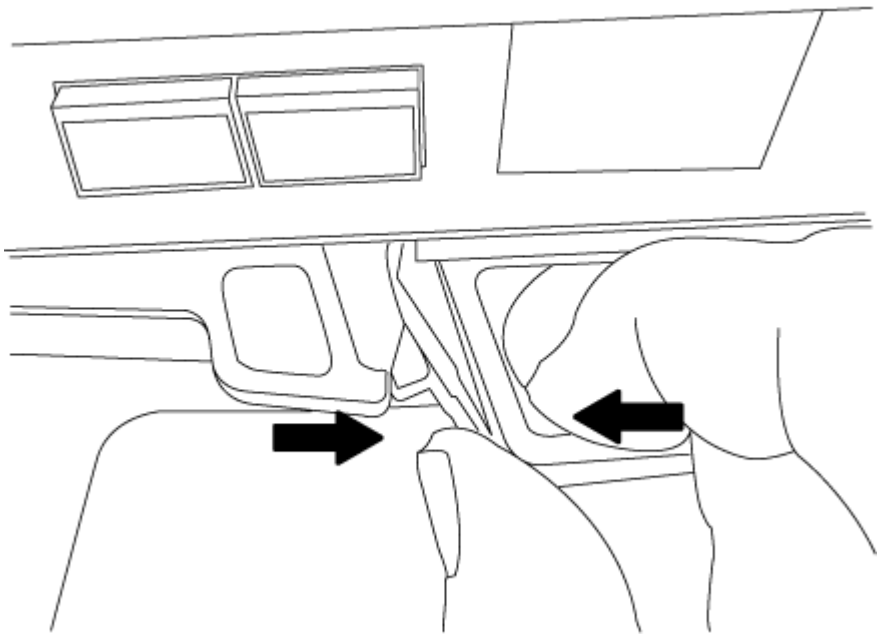

5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

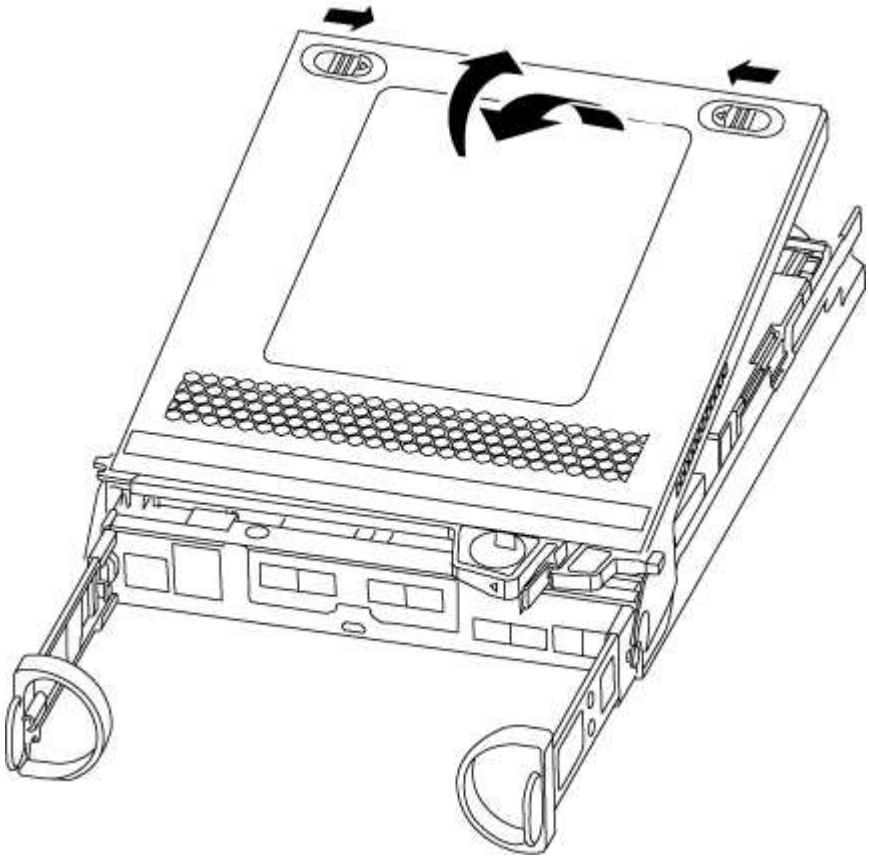

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

1. Locate the RTC battery.

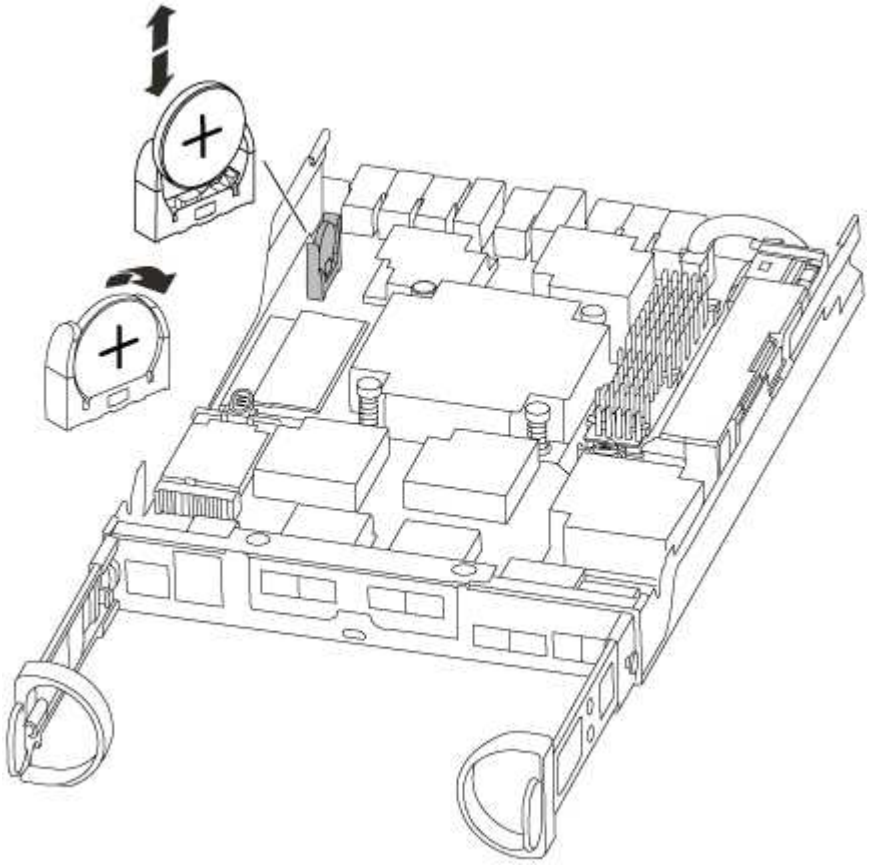

2. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

3. Remove the replacement battery from the antistatic shipping bag.
4. Locate the empty battery holder in the controller module.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and set time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Complete the replacement process

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

FAS500f systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Quick steps - FAS500f

This section gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Access the *Installation and Setup Instructions* PDF poster:

- English: FAS500f Installation and Setup Instructions
- Japanese: FAS500f Systems Installation and Setup Instructions
- Chinese: FAS500f Systems Installation and Setup Instructions

Video steps - FAS500f

The following video shows how to install and cable your new system.

Animation - Install and Setup of a FAS500f

Detailed steps - FAS500f

This section gives detailed step-by-step instructions for installing a FAS500f system.

Step 1: Prepare for installation

To install your FAS500f system, you need to create an account and register the system. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the NetApp Hardware Universe (HWU) for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

What you need

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register (NetApp Product Registration) your system.
4. Download and install NetApp Downloads: Config Advisor on your laptop.
5. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
25 GbE cable	X66240A-05 (112-00595), 0.5m; X66240-2 (112-00573), 2m		Cluster interconnect network
	X66240A-2 (112-00598), 2m; X66240A-5 (112-00600), 5m		Data
100 GbE cable	X66211-2 (112-00574), 2m; X66211-5 (112-00576), 5m		Storage
RJ-45 (order dependent)	Not applicable		Management network (BMC and wrench port) and Ethernet data (e0a and e0b)
Fibre Channel	X66250-2 (112-00342) 2m; X66250-5 (112-00344) 5m; X66250-15 (112-00346) 15m; X66250-30 (112-00347) 30m		

Type of cable…	Part number and length	Connector type	For…
Micro-USB console cable	Not applicable		Console connection during software setup
Power cables	Not applicable		Powering up the system

6. Review the ONTAP Configuration Guide and collect the required information listed in that guide.

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

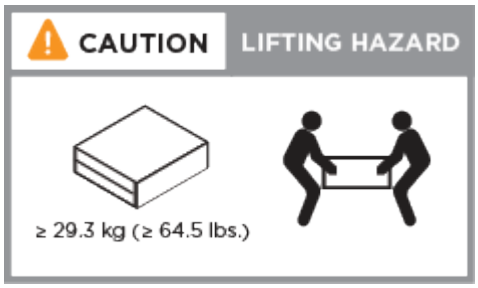

3. Identify and manage cables because this system does not have a cable management device.
4. Place the bezel on the front of the system.

Step 3: Cable controllers

There is required cabling for your platform's cluster using the two-node switchless cluster method or the cluster interconnect network method. There is optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cable to a host network and storage.

Required cabling: Cable controllers to a cluster

Cable the controllers to a cluster by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Cable a two-node switchless cluster

The management, Fibre Channel, and data or host network ports on the controller modules are connected to switches. The cluster interconnect ports are cabled on both controller modules.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

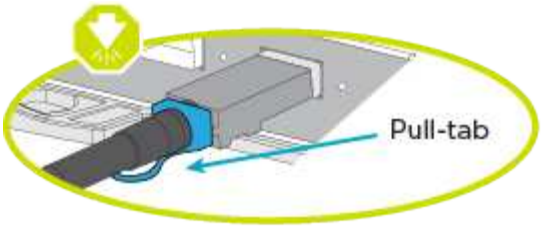

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to complete the cabling between the controllers and the switches:

Animation - Cable a two-node switchless cluster

Step	Perform on each controller
1	Cable the cluster interconnect ports to each other with the 25GbE cluster interconnect cable  : • e0c to e0c • e0d to e0d 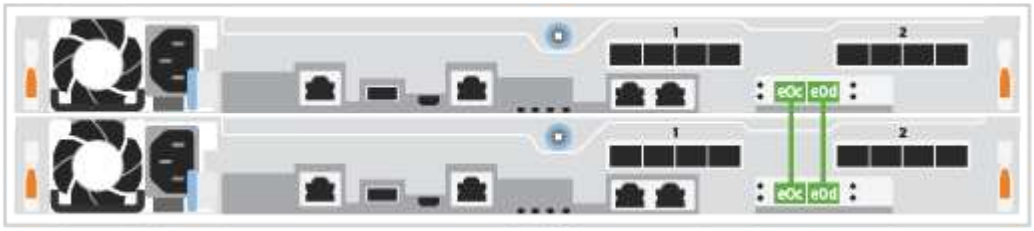
2	Cable the wrench ports to the management network switches with the RJ45 cables. 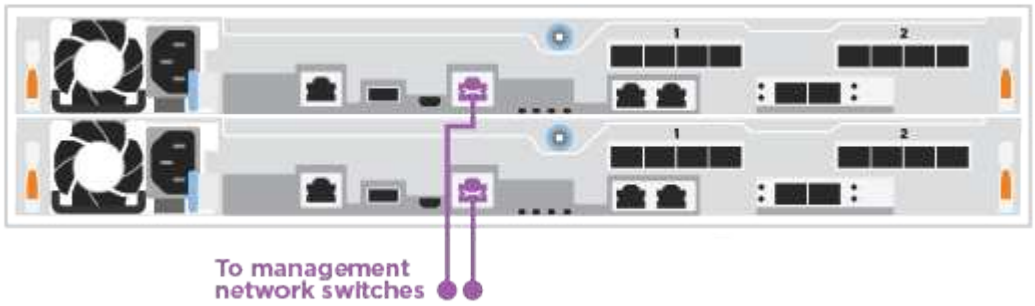
	DO NOT plug in the power cords at this point.

To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 2: Cable a switched cluster

All ports on the controllers are connected to switches; cluster interconnect, management, Fibre Channel, and data or host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

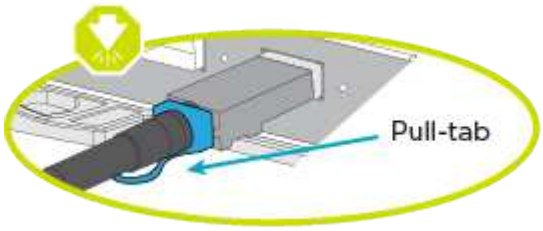

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to complete the cabling between the controllers and the switches:

Animation - Cable a switched cluster

Step	Perform on each controller
1	Cable the cluster interconnect ports to the 25 GbE cluster interconnect switches. • e0c • e0d 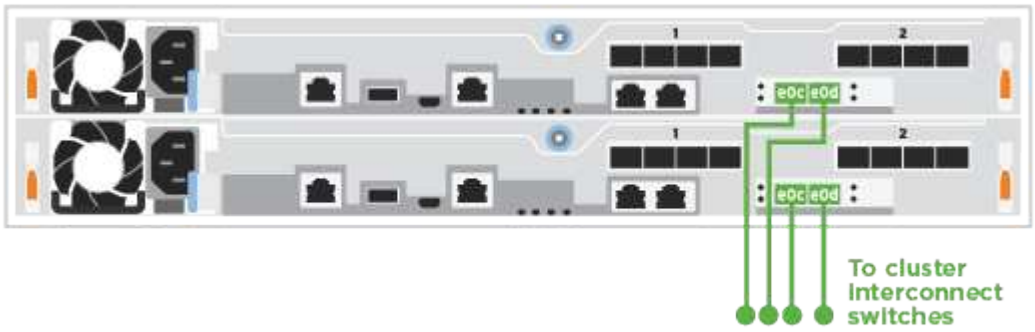

Step	Perform on each controller
2	Cable the wrench ports to the management network switches with the RJ45 cables. 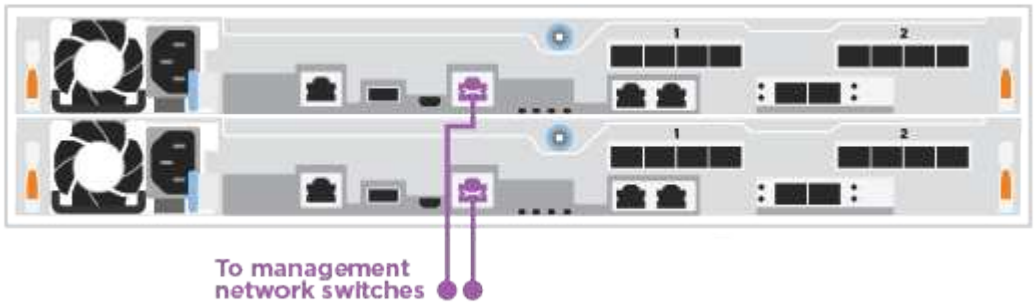
	DO NOT plug in the power cords at this point.

To complete setting up your system, see Step 4: Complete system setup and configuration.

Optional cabling: Cable configuration-dependent options

You have configuration-dependent optional cabling to the Fibre Channel or iSCSI host networks or direct-attached storage. This cabling is not exclusive; you can have cabling to a host network and storage.

Option 1: Cable to a Fibre Channel host network

Fibre Channel ports on the controllers are connected to Fibre Channel host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

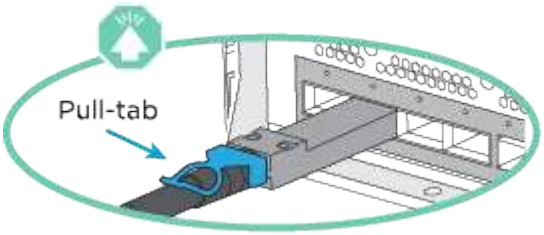

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Step	Perform on each controller module
1	Cable ports 2a through 2d to the FC host switches. 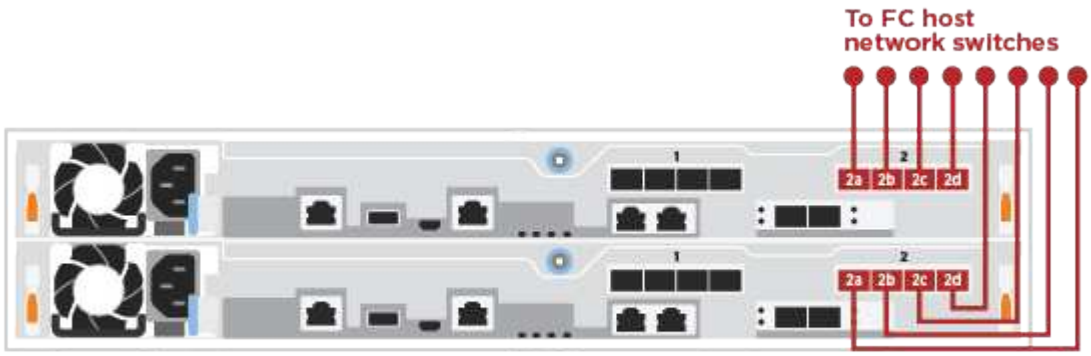
2	To perform other optional cabling, choose from: • Option 2: Cable to a 25GbE data or host network • Option 3: Cable the controllers to a single drive shelf
3	To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 2: Cable to a 25GbE data or host network

25GbE ports on the controllers are connected to 25GbE data or host network switches.

Before you begin

Contact your network administrator for information about connecting the system to the switches.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

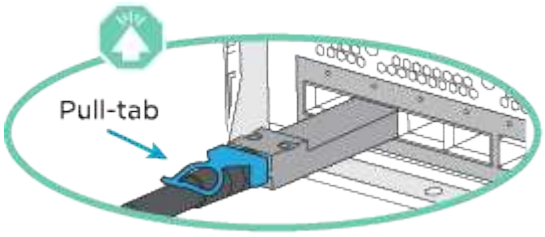

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Step	Perform on each controller module
1	Cable ports e4a through e4d to the 10GbE host network switches. 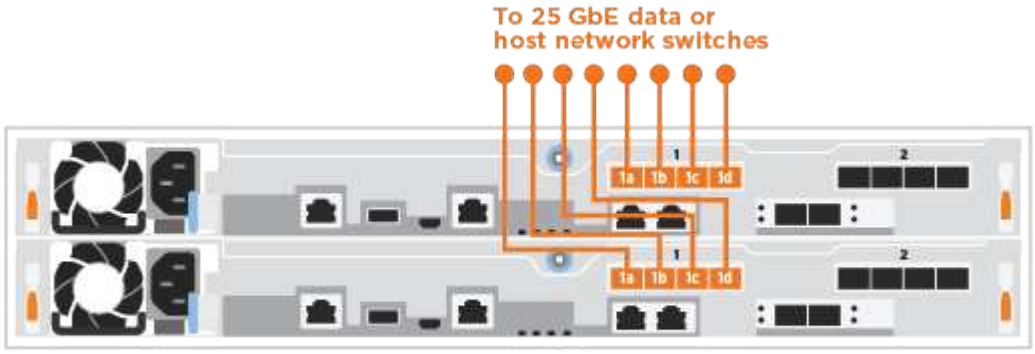
2	To perform other optional cabling, choose from: • Option 1: Cable to a Fibre Channel host network • Option 3: Cable the controllers to a single drive shelf
3	To complete setting up your system, see Step 4: Complete system setup and configuration.

Option 3: Cable the controllers to a single drive shelf

You must cable each controller to the NSM modules on the NS224 drive shelf.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

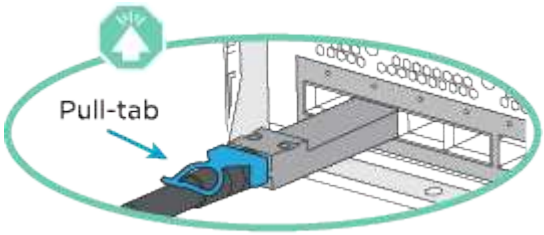

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Use the animation or the tabulated steps to complete the cabling between the controllers and the single shelf:

Animation - Cable the controllers to a single NS224

Step	Perform on each controller module
1	Cable controller A to the shelf: 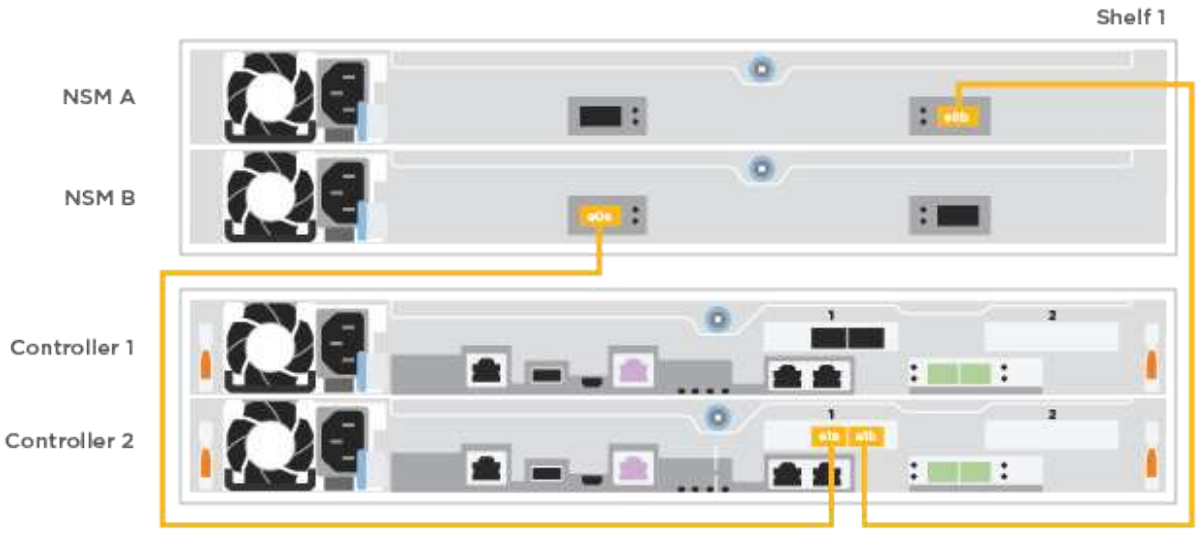
2	Cable controller B to the shelf: 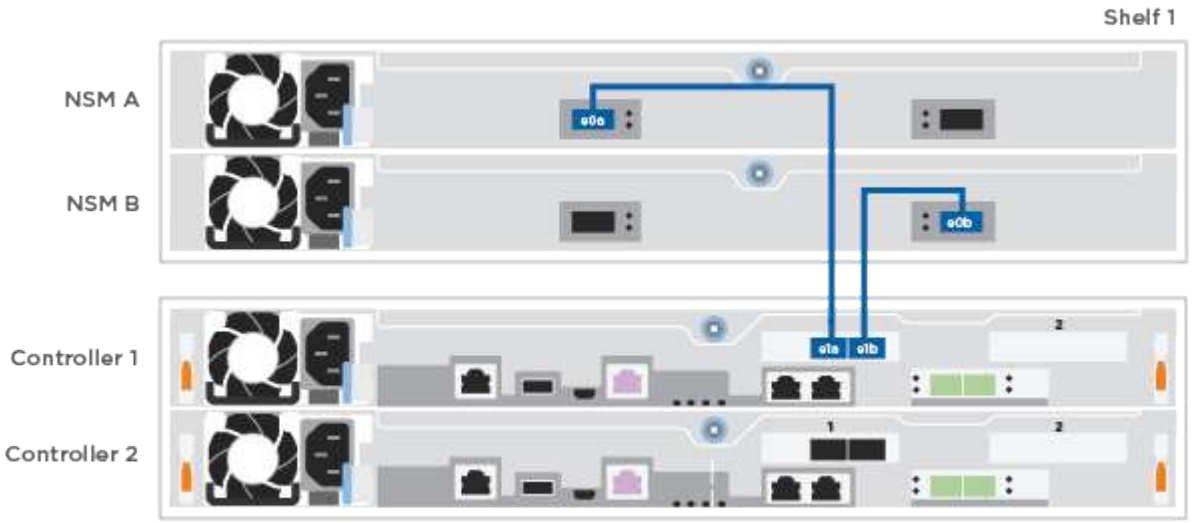

To complete setting up your system, see Step 4: Complete system setup and configuration.

Step 4: Complete system setup and configuration

Complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Complete system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
2. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

3. Use the animation to connect your laptop to the Management switch:

   Animation - Connect your laptop to the Management switch

4. Select an ONTAP icon listed to discover:

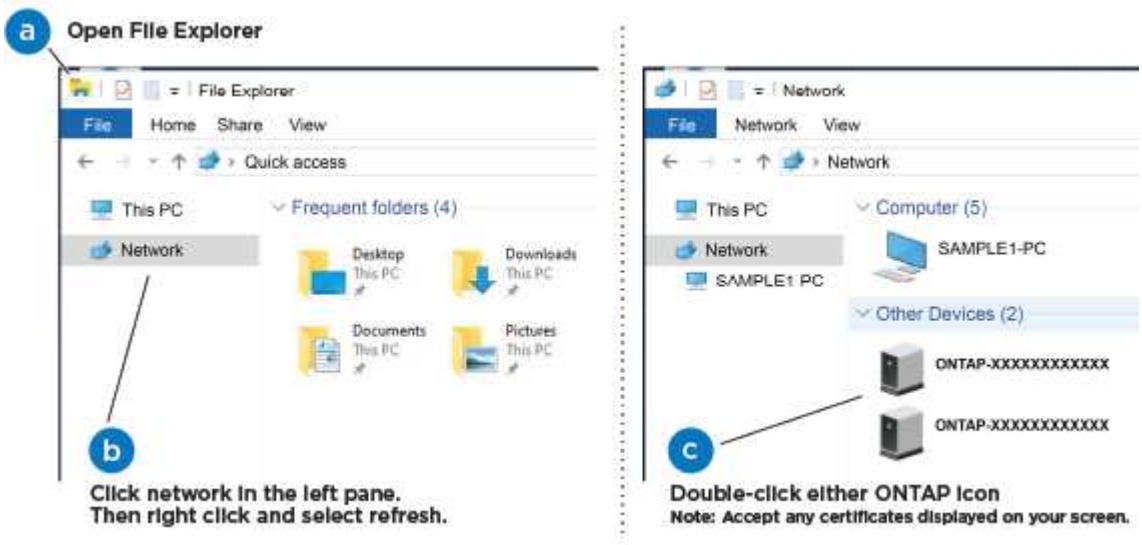

   a. Open File Explorer.
   b. Click **Network** in the left pane.
   c. Right-click and select **refresh**.
   d. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

      System Manager opens.

5. Use System Manager guided setup to configure your system using the data you collected in the ONTAP Configuration Guide.
6. Verify the health of your system by running Config Advisor.
7. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Complete system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:

a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

> See your laptop or console's online help for how to configure the console port.

b. Connect the laptop or console to the switch on the management subnet.

c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

4. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

   > The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the ONTAP Configuration Guide.
5. Verify the health of your system by running Config Advisor.
6. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain FAS500f hardware

For the FAS500f storage system, you can perform maintenance procedures on the

following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

Fan

The fan cools the controller.

Mezzanine card

A Mezzanine card is a printed circuit board that plugs directly into another plug-in card.

NVMEM battery

A battery is included with the controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real-time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - FAS500f

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots.

You must have a USB flash drive, formatted to MBR/FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct controller:
  - The *impaired* controller is the controller on which you are performing maintenance.
  - The *healthy* controller is the HA partner of the impaired controller.

Check encryption key support and status - FAS500f

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.

ONTAP version	Run this command
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the controller - FAS500f

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - FAS500f

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

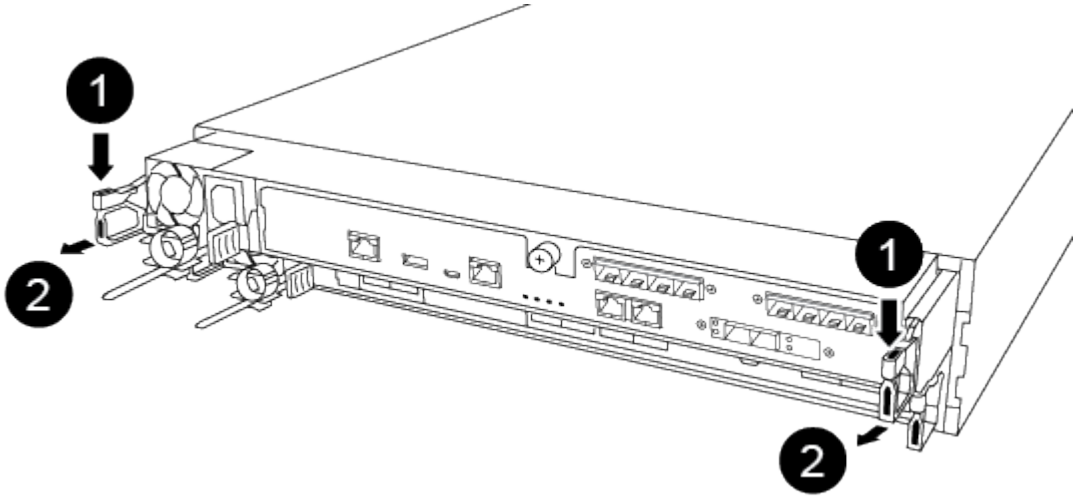

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

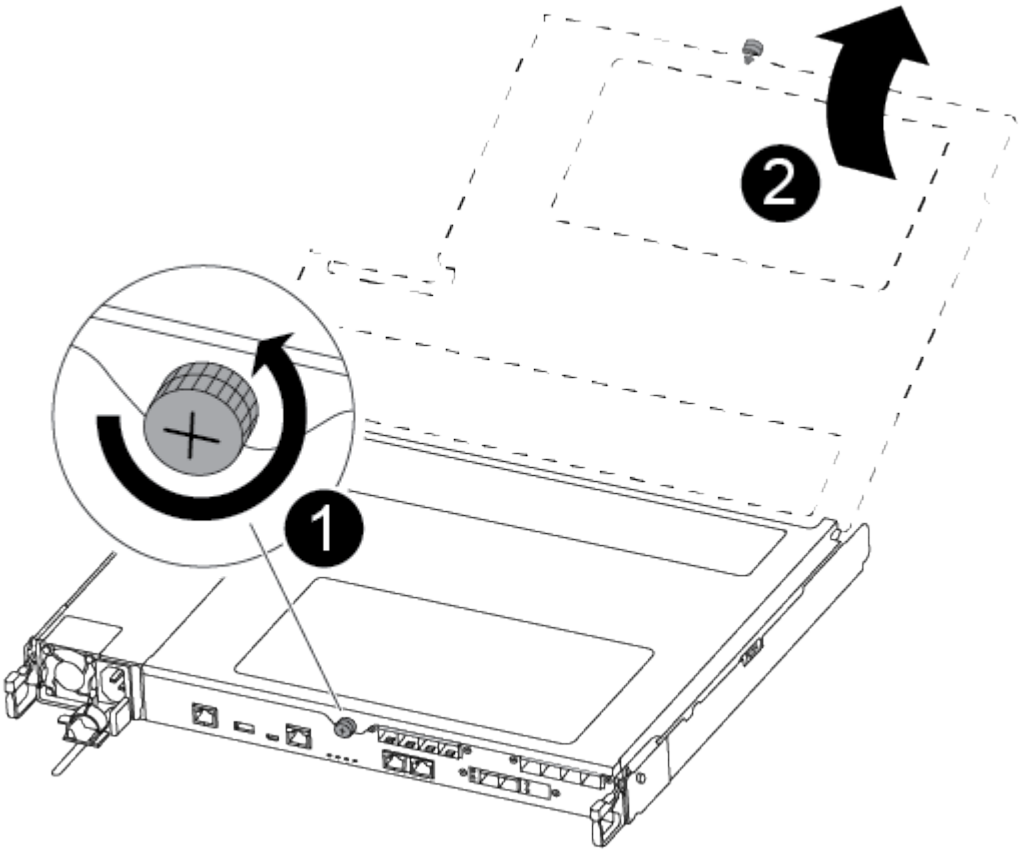

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

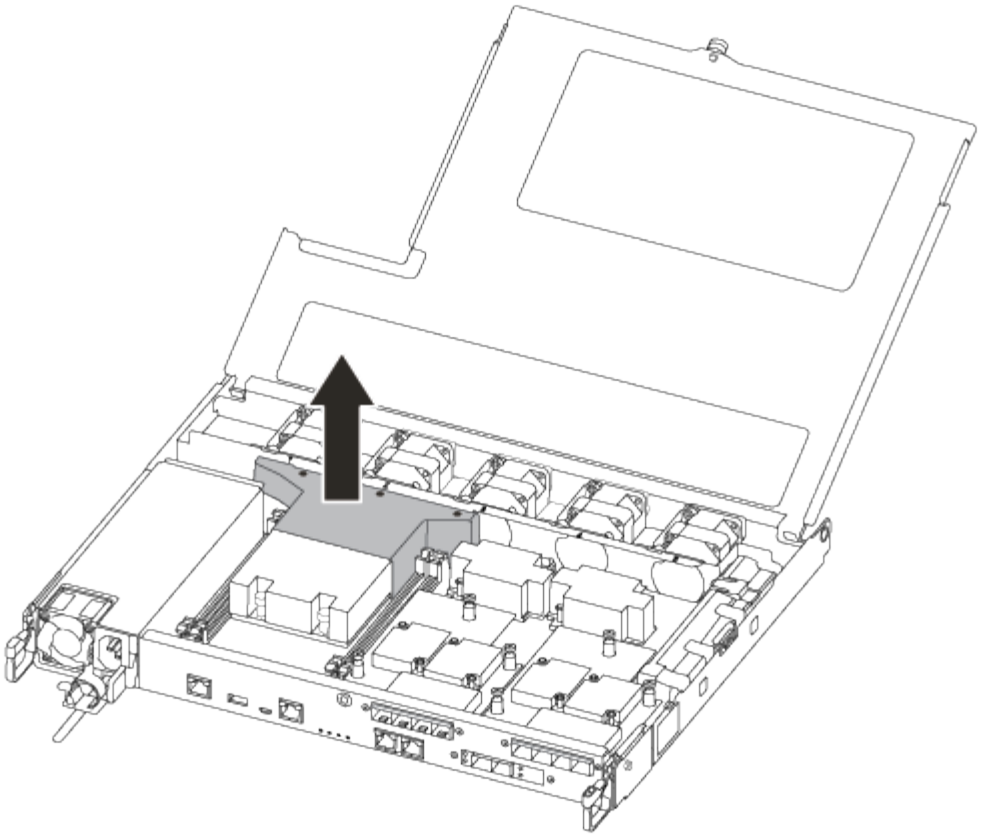

Step 2: Replace the boot media

You locate the failed boot media in the controller module by removing the air duct on the controller module before you can replace the boot media.

You need a #1 magnetic Phillips head screwdriver to remove the screw that holds the boot media in place. Due to the space constraints within the controller module, you should also have a magnet to transfer the screw on to so that you do not loose it.

You can use the following video or the tabulated steps to replace the boot media:

Animation - Replace the boot media

1. Locate and replace the impaired boot media from the controller module.

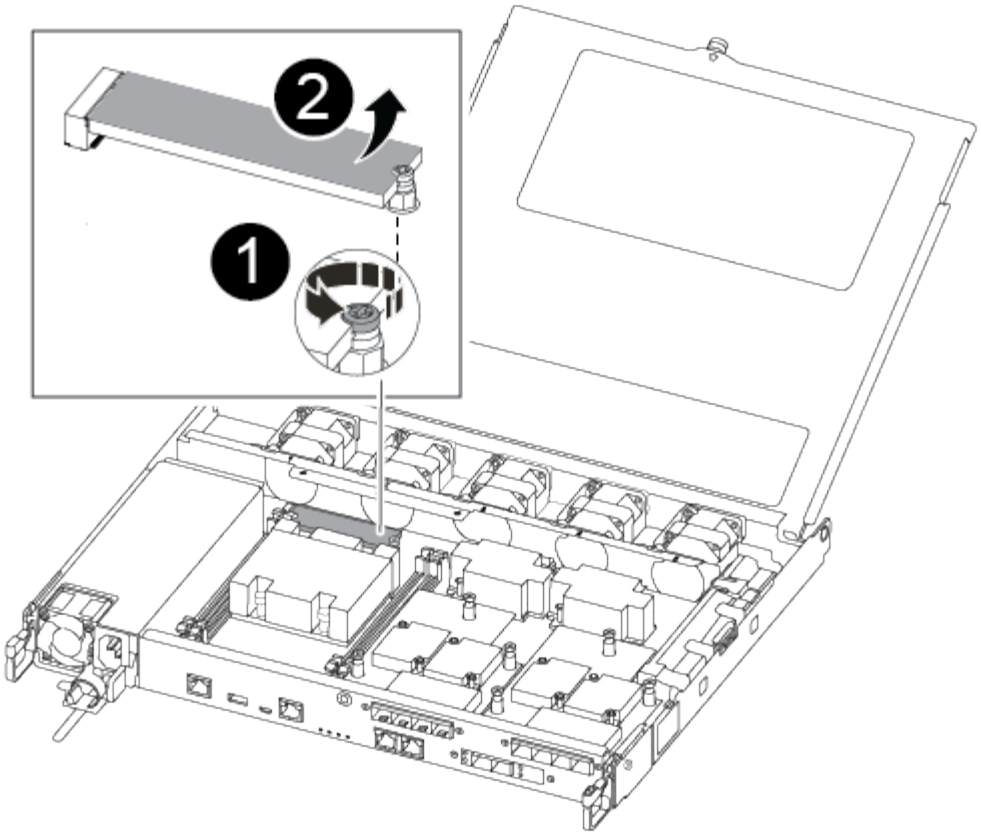

1	Remove the screw securing the boot media to the motherboard in the controller module.
2	Lift the boot media out of the controller module.

a. Using the #1 magnetic screwdriver, remove the screw from the impaired boot media, and set it aside safely on the magnet.

b. Gently lift the impaired boot media directly out of the socket and set it aside.

c. Remove the replacement boot media from the antistatic shipping bag and align it into place on the controller module.

d. Using the #1 magnetic screwdriver, insert and tighten the screw on the boot media.

Do not apply force when tightening the screw on the boot media; you might crack it.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed is without a boot image so you need to transfer a boot image using a USB flash drive.

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the **Downloads** section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download

button.

- If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
2. Download the service image to your work space on your laptop.
3. Unzip the service image.

   If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

   There are two folders in the unzipped service image file:

   - `boot`
   - `efi`

4. Copy the `efi` folder to the top directory on the USB flash drive.

   The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

5. Remove the USB flash drive from your laptop.
6. If you have not already done so, install the air duct.

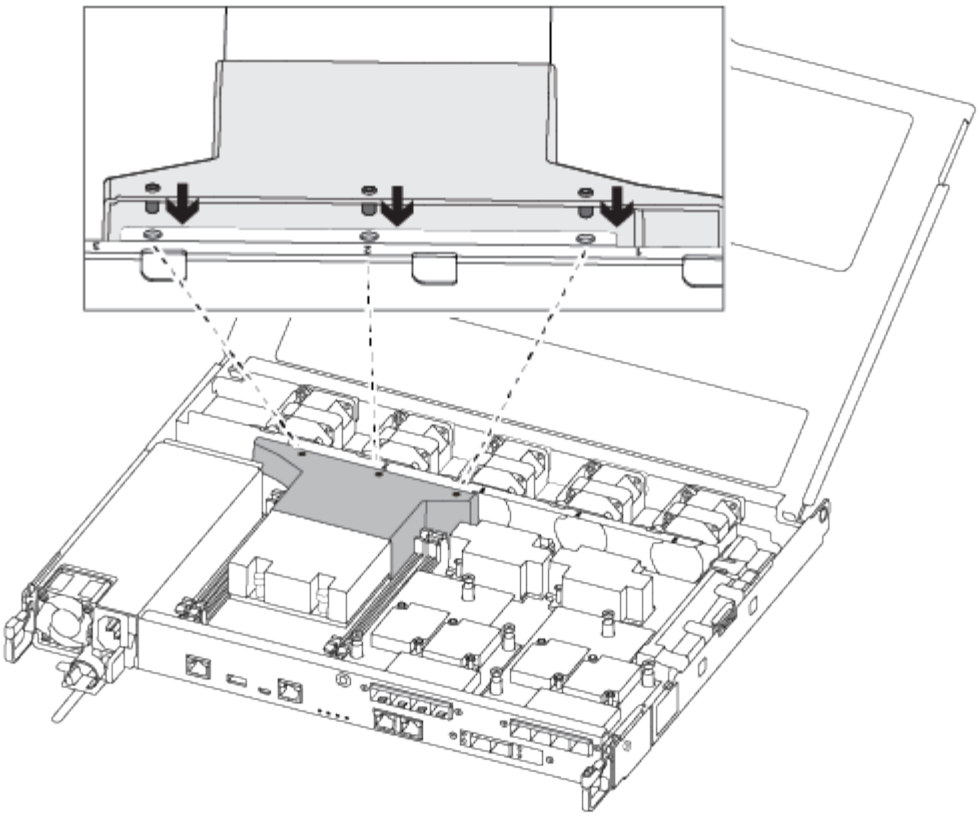

7. Close the controller module cover and tighten the thumbscrew.

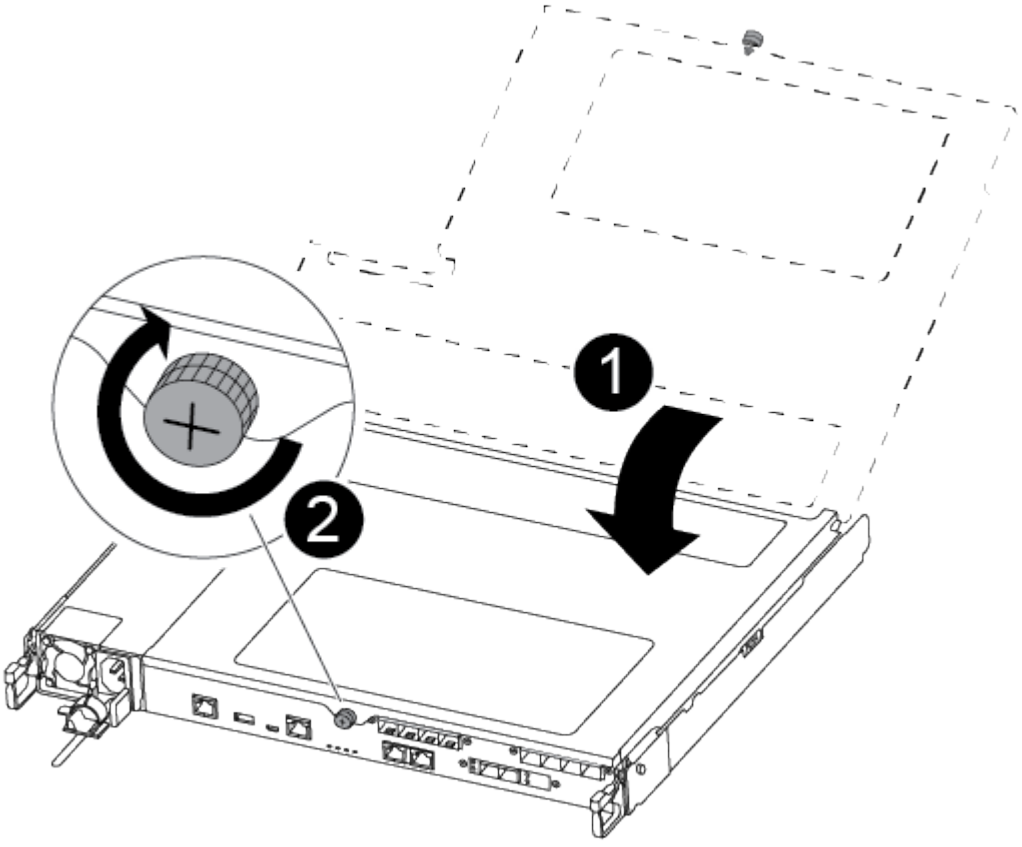

1	Controller module cover
2	Thumbscrew

8. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
9. Plug the power cable into the power supply and reinstall the power cable retainer.
10. Insert the USB flash drive into the USB slot on the controller module.

    Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

11. Push the controller module all the way into the chassis:
12. Place your index fingers through the finger holes from the inside of the latching mechanism.
13. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
14. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

    The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

    The controller module should be fully inserted and flush with the edges of the chassis.

Boot the recovery image - FAS500f

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive:

   ```
   boot_recovery
   ```

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

If your system has…	Then…

No network connection and is in a MetroCluster IP configuration	a. Press `n` when prompted to restore the backup configuration.
	b. Reboot the system when prompted by the system.
	c. Wait for the iSCSI storage connections to connect.
	You can proceed after you see the following messages:
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_auxiliary, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_partner, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_auxiliary, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_partner, address: ip-address).`
	d. Select the **Update flash from backup config** (sync flash) option from the displayed menu.
	If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:
   - If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed

- If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.
8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore encryption - FAS500f

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************************
**************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*********************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*********************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - FAS500f

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - FAS500f

To replace the chassis, you must move the bezel, controller modules, and NVMe drives from the impaired chassis to the replacement chassis, and then remove the impaired chassis from the equipment rack or system cabinet and install the replacement chassis in its place.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - FAS500f

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - FAS500f

Move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

You can use the following video or the tabulated steps to replace the chassis; it assumes the removal and replacement of the bezel:

Animation - Replace the chassis

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

   > If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

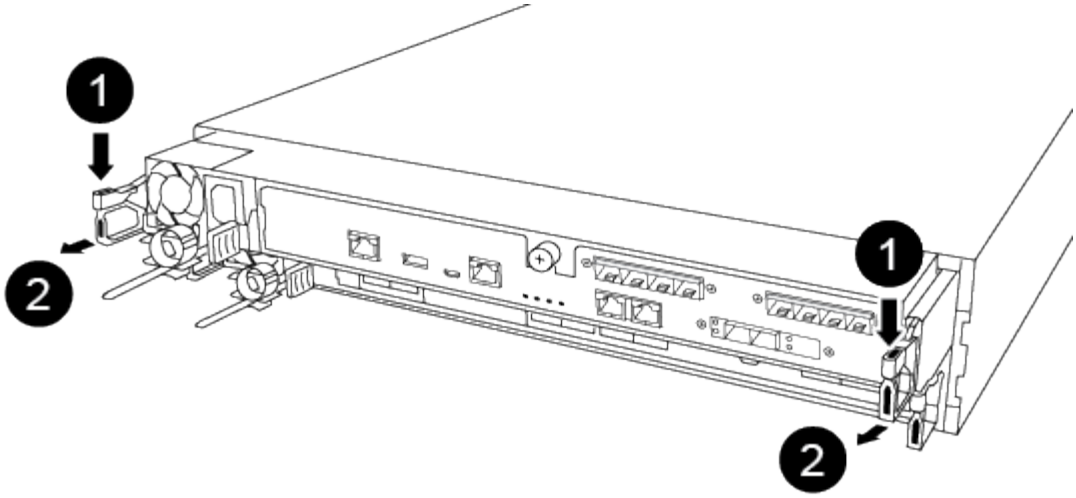

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      When removing a drive, always use two hands to support its weight.

      Drives are fragile. Handle them as little as possible to prevent damage to them.
3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.
5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up

and against the drive holder.

Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Plug the power cables into the power supplies and reinstall the power cable retainers.
4. Insert the controller module into the chassis:

   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to

interrupt the boot process.

The controller module should be fully inserted and flush with the edges of the chassis.

5. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - FAS500f

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis *HA-state*`

      The value for HA-state can be one of the following:

      - `ha`
      - `mcc`
      - `mccip`
      - `non-ha`

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. Reinstall the bezel on the front of the system.

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - FAS500f

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - FAS500f

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - FAS500f

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

You can use the following video or the tabulated steps to replace a controller module:

Animation - Replace a controller module

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

> If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

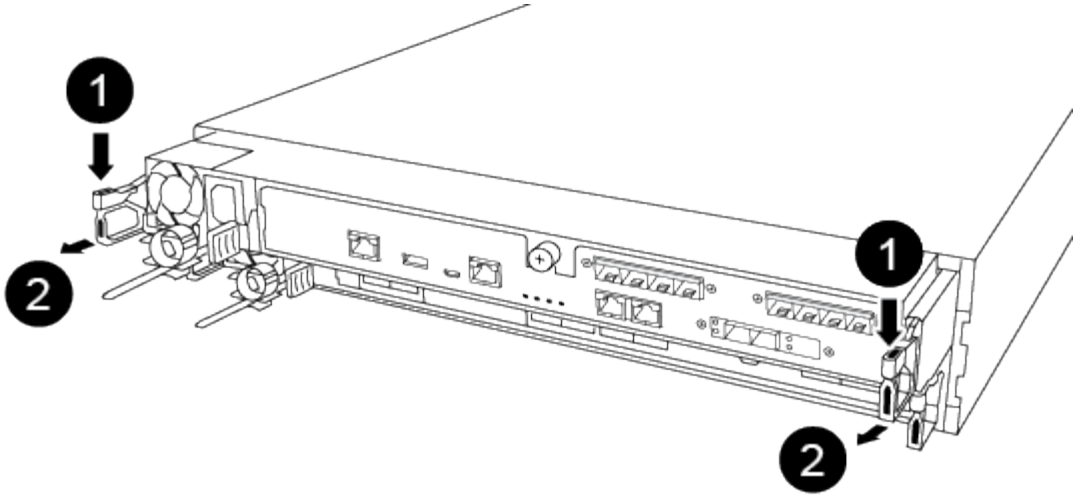

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

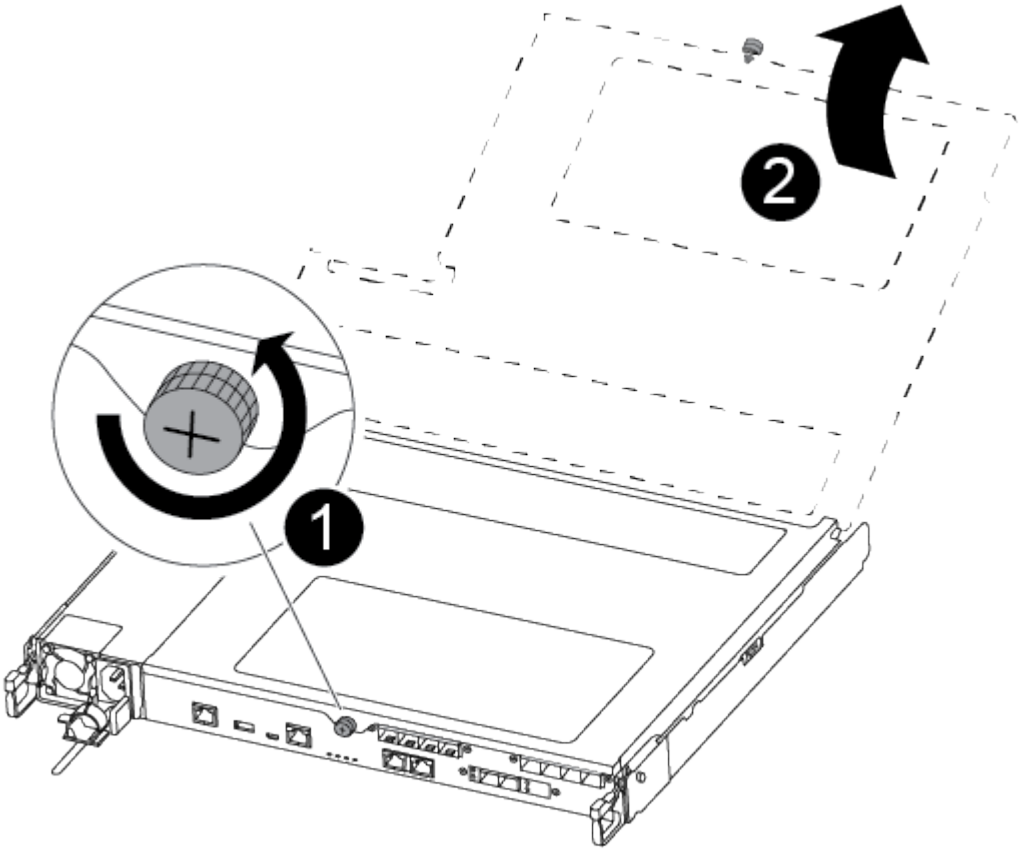

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

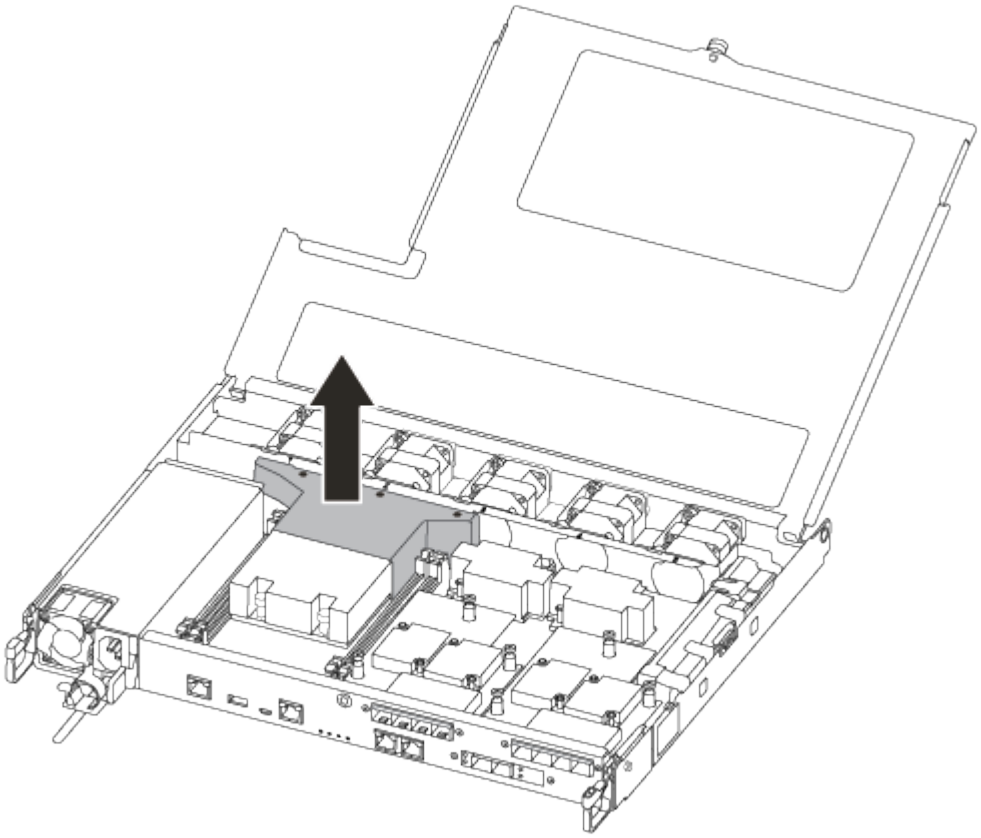

Step 2: Move the power supply

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

1. Disconnect the power supply.
2. Open the power cable retainer, and then unplug the power cable from the power supply.
3. Unplug the power cable from the power source.
4. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

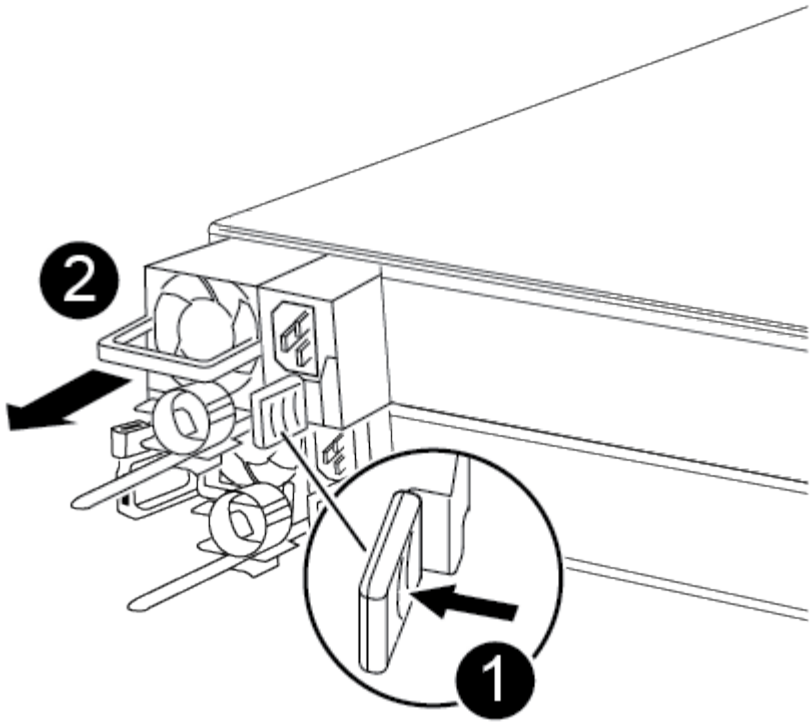

1	Blue power supply locking tab
2	Power supply

5. Move the power supply to the new controller module, and then install it.
6. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

You must move the fans from the impaired controller module to the replacement module when replacing a failed controller module.

1. Remove the fan module by pinching the side of the fan module, and then lifting the fan module straight out of the controller module.

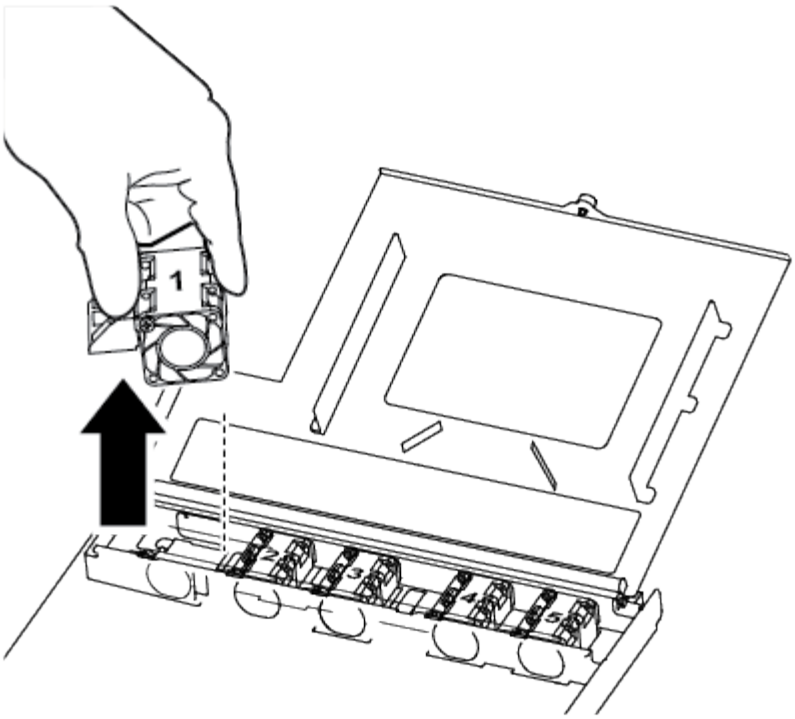

1	Fan module

2. Move the fan module to the replacement controller module, and align the edges of the fan module with the opening in the controller module, and then slide the fan module in.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the boot media

There is one boot media device in the AFF A250 under the air duct in the controller module. You must move it from the impaired controller module to the replacement controller module.

You need a #1 magnetic Phillips head screwdriver to remove the screw that holds the boot media in place. Due to the space constraints within the controller module, you should also have a magnet to transfer the screw on to so that you do not lose it.

1. Locate and move the boot media from the impaired controller module to the replacement controller module.

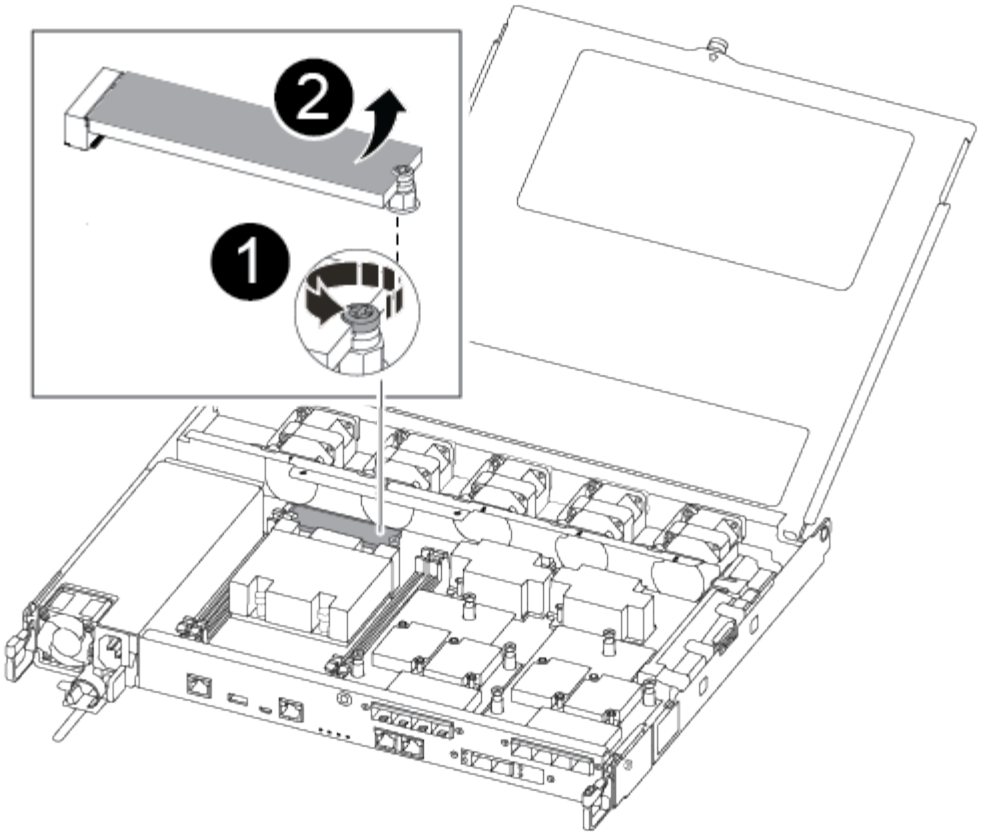

1	Remove the screw securing the boot media to the motherboard in the impaired controller module.
2	Lift the boot media out of the impaired controller module.

a. Using the #1 magnetic screwdriver, remove the screw from the boot media, and set it aside safely on the magnet.

b. Gently lift the boot media directly out of the socket and align it into place in the replacement controller module.

c. Using the #1 magnetic screwdriver, insert and tighten the screw on the boot media.

> Do not apply force when tightening the screw on the boot media; you might crack it.

Step 5: Move the DIMMs

To move the DIMMs, locate and move them from the impaired controller into the replacement controller and follow the specific sequence of steps.

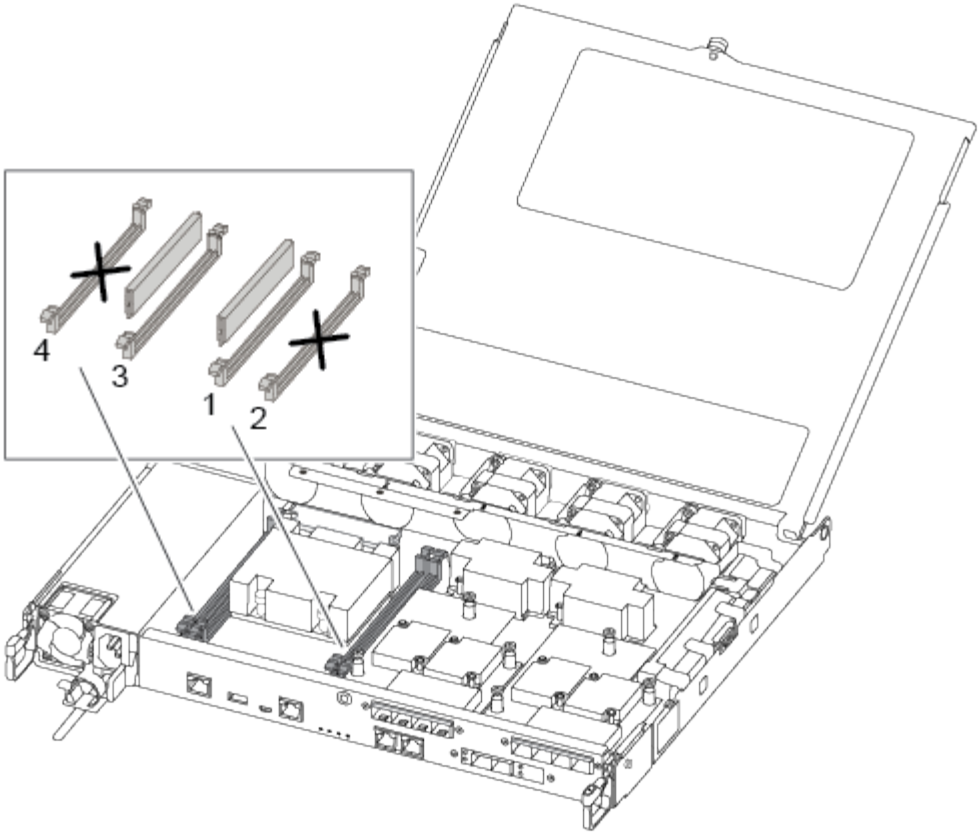

Install each DIMM into the same slot it occupied in the impaired controller module.

1. Slowly push apart the DIMM ejector tabs on either side of the DIMM, and slide the DIMM out of the slot.

   Hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

2. Locate the corresponding DIMM slot on the replacement controller module.
3. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the DIMM squarely into the socket.

   The DIMMs fit tightly in the socket. If not, reinsert the DIMM to realign it with the socket.

4. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.
5. Repeat these steps for the remaining DIMM.

Step 6: Move a mezzanine card

To move a mezzanine card, you must remove the cabling and any QSFPs and SFPs from the ports, move the mezzanine card to the replacement controller, reinstall any QSFPs and SFPs onto the ports, and cable the ports.

1. Locate and move the mezzanine cards from your impaired controller module.

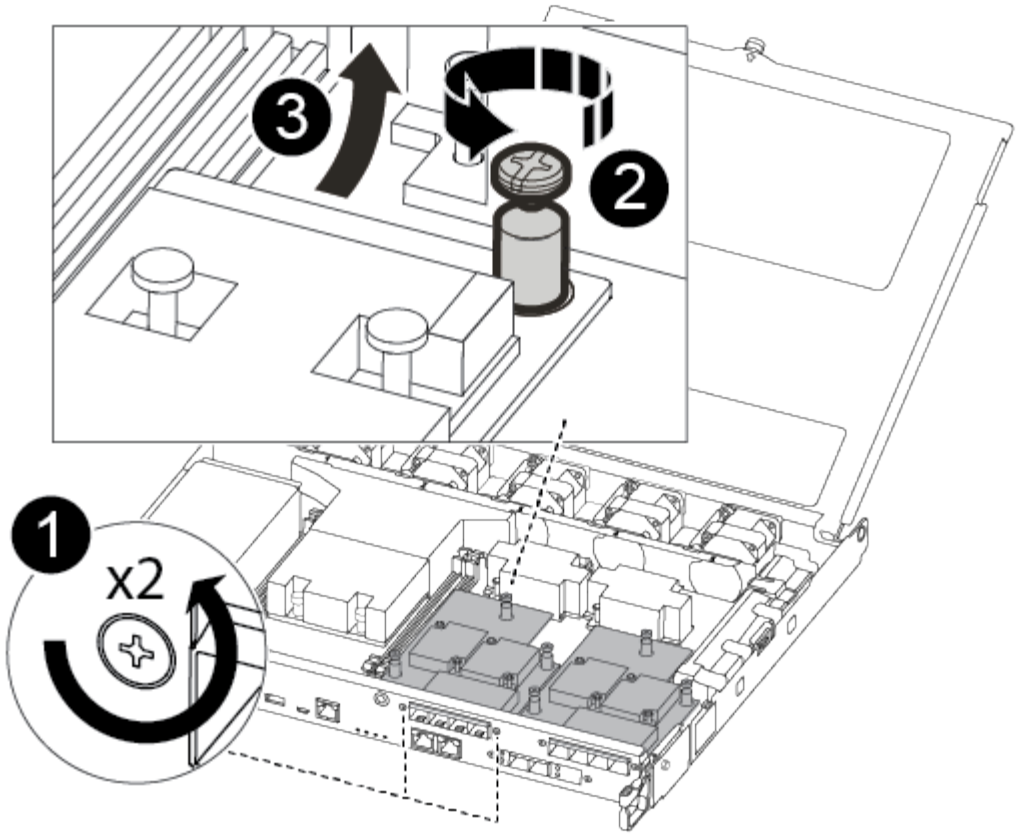

1	Remove screws on the face of the controller module.
2	Loosen the screw in the controller module.
3	Move the mezzanine card.

2. Unplug any cabling associated with the mezzanine card.

   Make sure that you label the cables so that you know where they came from.

   a. Remove any SFP or QSFP modules that might be in the mezzanine card and set it aside.
   b. Using the #1 magnetic screwdriver, remove the screws from the face of the impaired controller module and from the mezzanine card, and set them aside safely on the magnet.
   c. Gently lift the mezzanine card out of the socket and move it to the same position in the replacement controller.
   d. Gently align the mezzanine card into place in the replacement controller.
   e. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the replacement controller module and on the mezzanine card.

      Do not apply force when tightening the screw on the mezzanine card; you might crack it.

3. Repeat these steps if there is another mezzanine card in the impaired controller module.

4. Insert the SFP or QSFP modules that were removed onto the mezzanine card.

Step 7: Move the NV battery

When replacing the controller module, you must move the NV battery from the impaired controller module to the replacement controller module.

1. Locate and move the NVMEM battery from your impaired controller module to the replacement controller module.

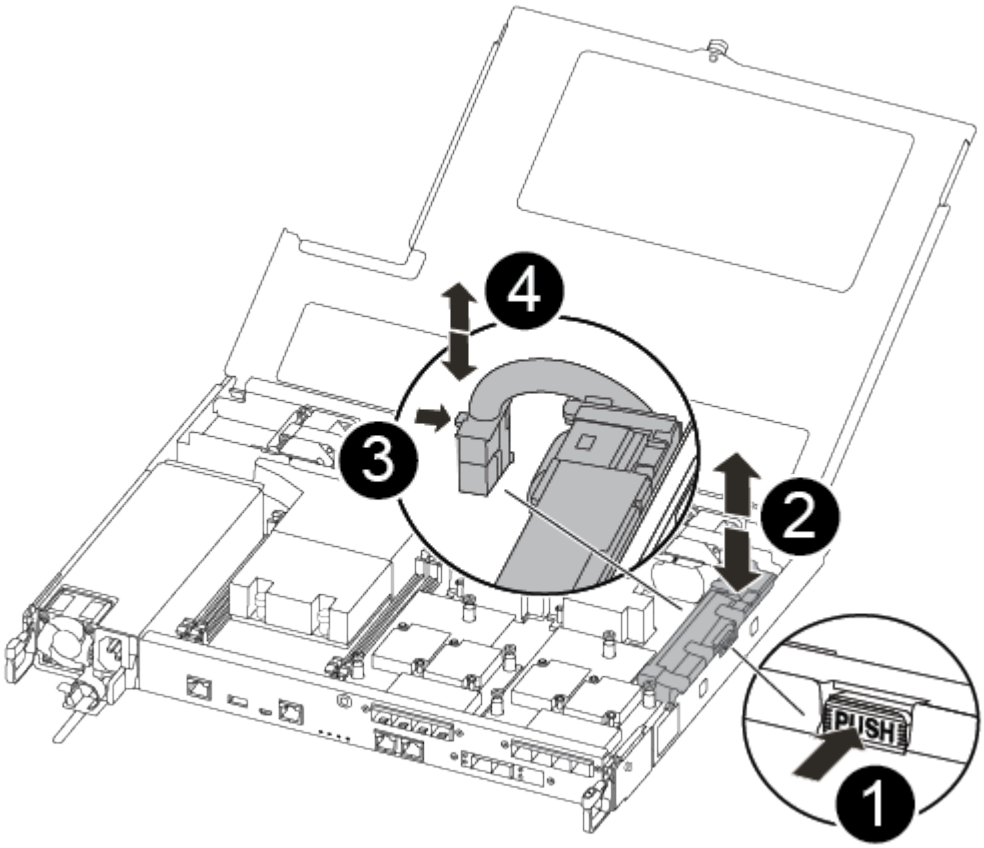

1	Squeeze the clip on the face of the battery plug.
2	Unplug the battery cable from the socket.
3	Grasp the battery and press the blue locking tab marked PUSH.
4	Lift the battery out of the holder and controller module.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.

4. Locate the corresponding NV battery holder on the replacement controller module and align the NV battery to the battery holder.
5. Insert the NV battery plug into the socket.
6. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
7. Press firmly down on the battery pack to make sure that it is locked into place.

Step 8: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

You can use the following illustration or the written steps to install the replacement controller module in the chassis.

1. If you have not already done so, install the air duct.

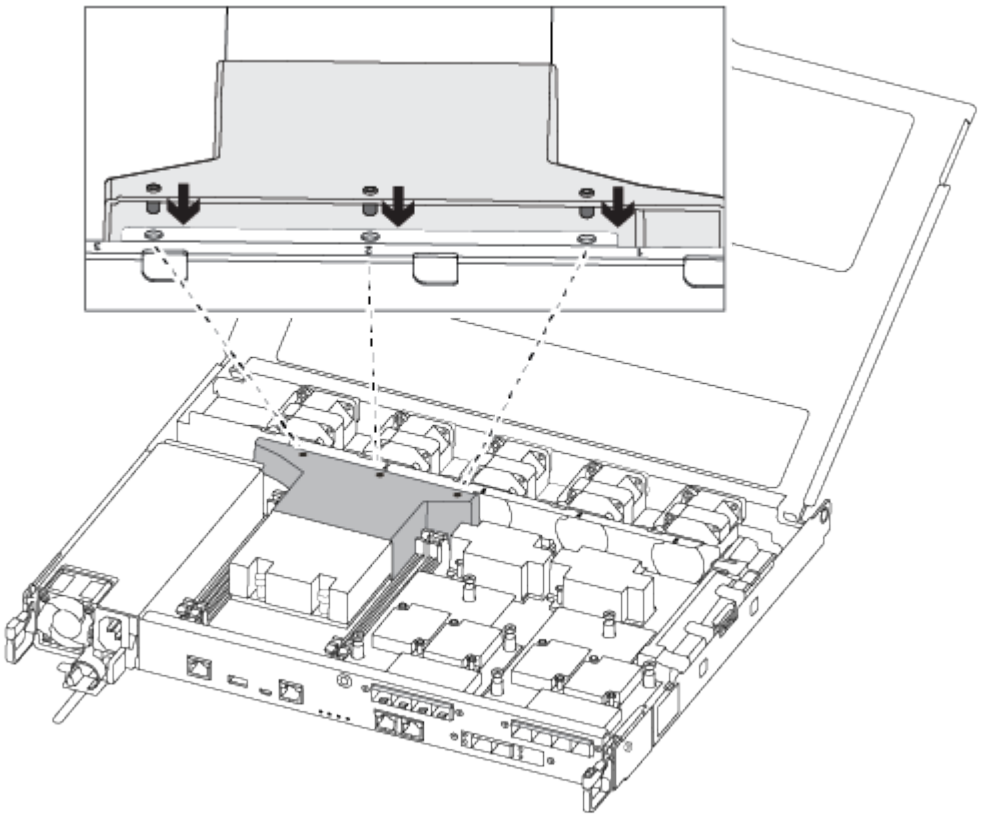

2. Close the controller module cover and tighten the thumbscrew.

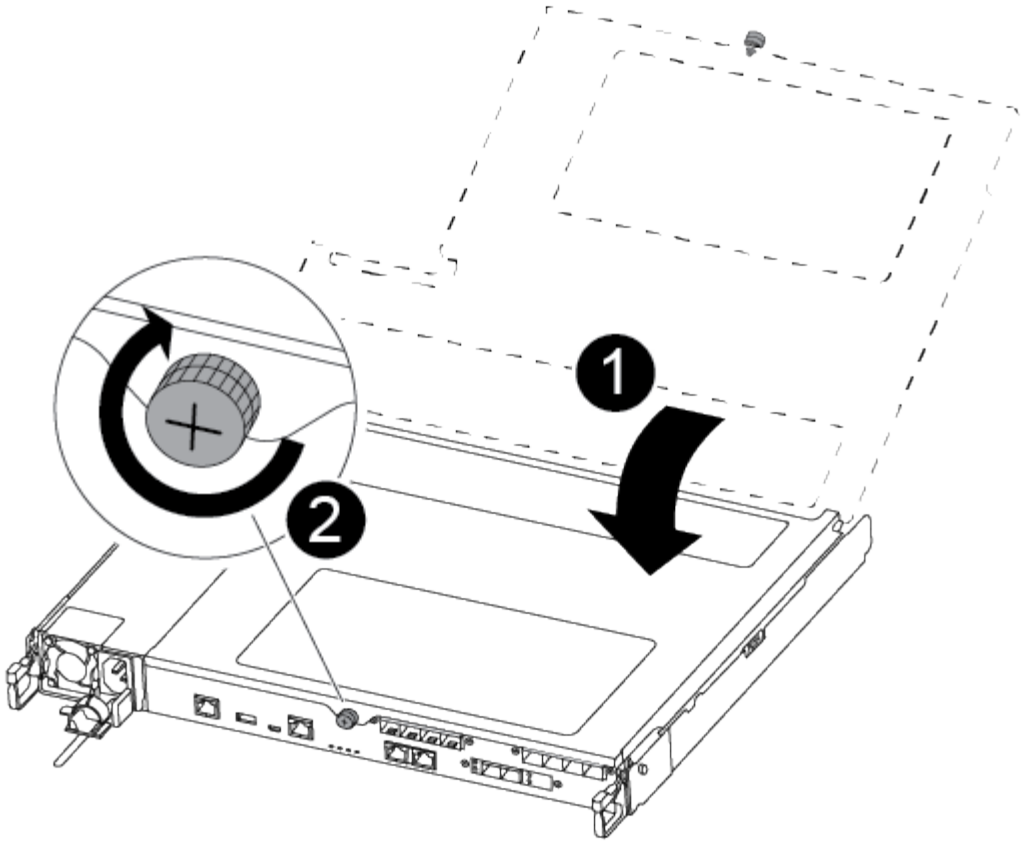

1	Controller module cover
2	Thumbscrew

3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Insert the controller module into the chassis.
6. Ensure the latching mechanism arms are locked in the fully extended position.
7. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
8. Place your index fingers through the finger holes from the inside of the latching mechanism.
9. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
10. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching

mechanisms snap into place.

The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

The controller module should be fully inserted and flush with the edges of the chassis.

Restore and verify the system configuration - FAS500f

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the chassis

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - ha
   - mcc
   - mccip
   - non-ha

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - FAS500f

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

   The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The *replacement* controller is the current owner of the disks on the disaster site.

     Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - FAS500f

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS500f

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

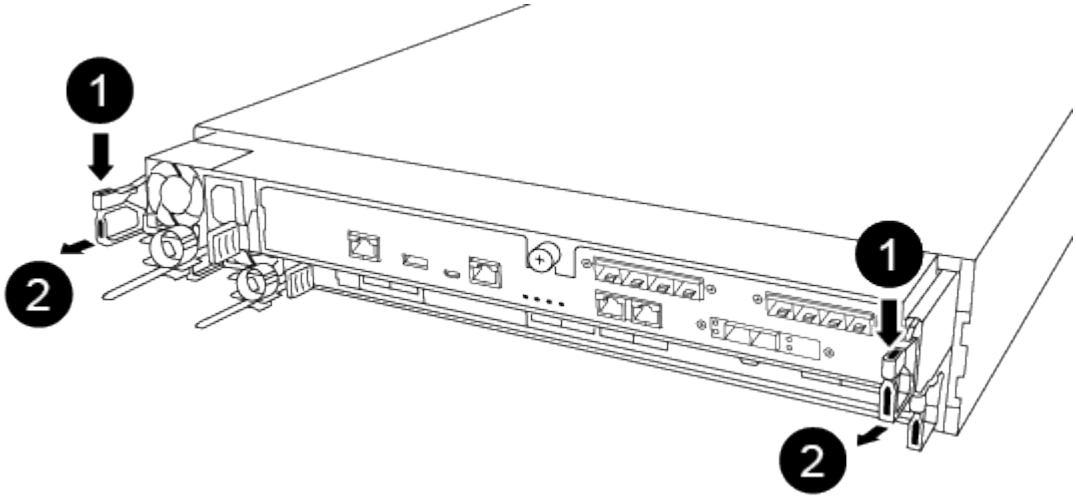

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

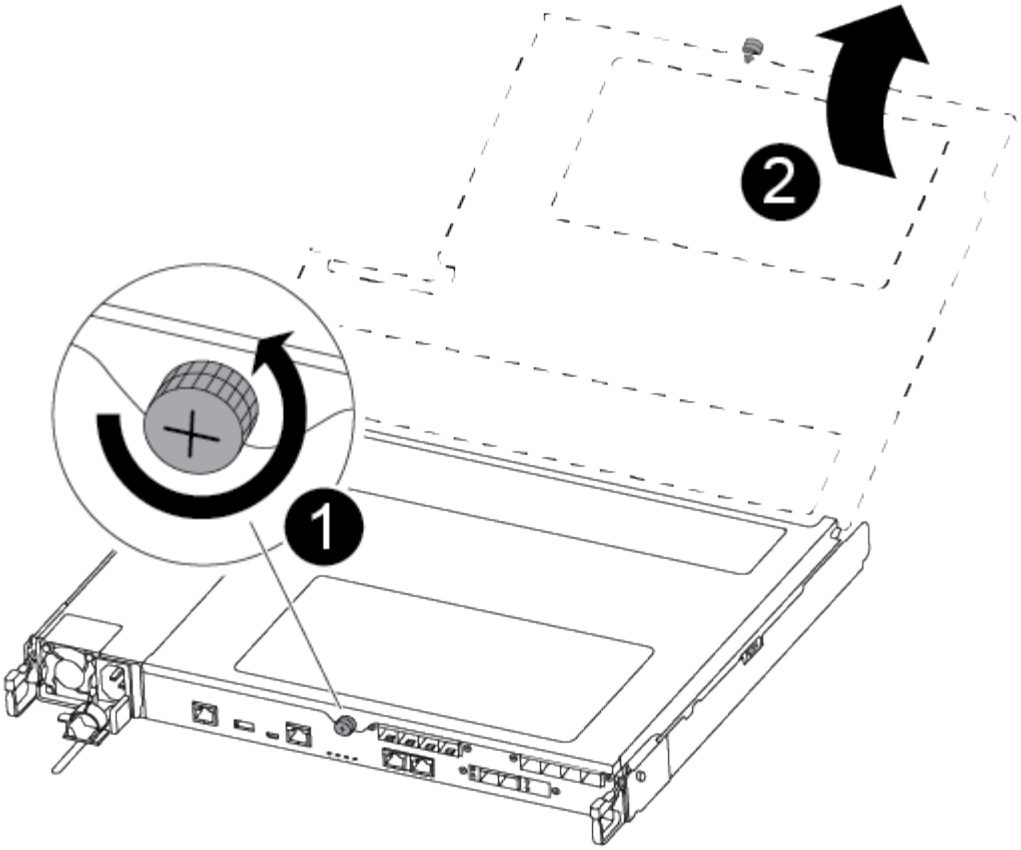

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

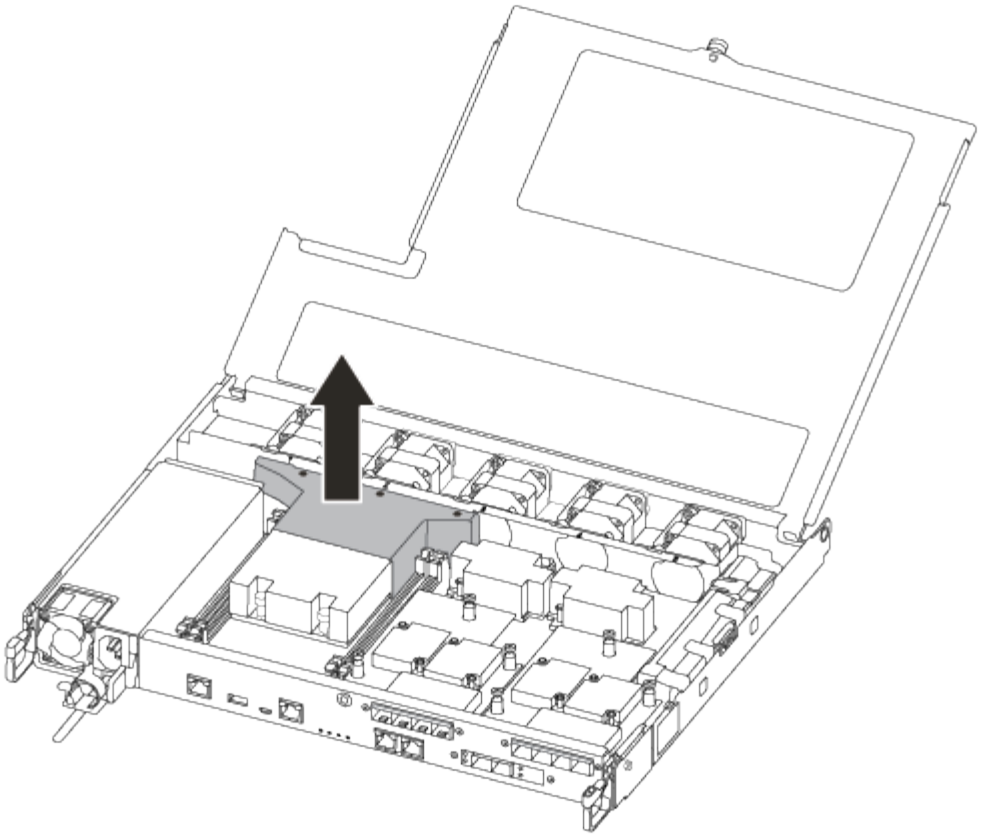

Step 3: Replace a DIMM

To replace a DIMM, you must locate it in the controller module using the DIMM map label on top of the air duct and then replace it following the specific sequence of steps.

You can use the following video or the tabulated steps to replace a DIMM:

Animation - Replace a DIMM

1. Replace the impaired DIMM on your controller module.

   The DIMMs are in slot 3 or 1 on the motherboard. Slot 2 and 4 are left empty. Do not attempt to install DIMMs into these slots.

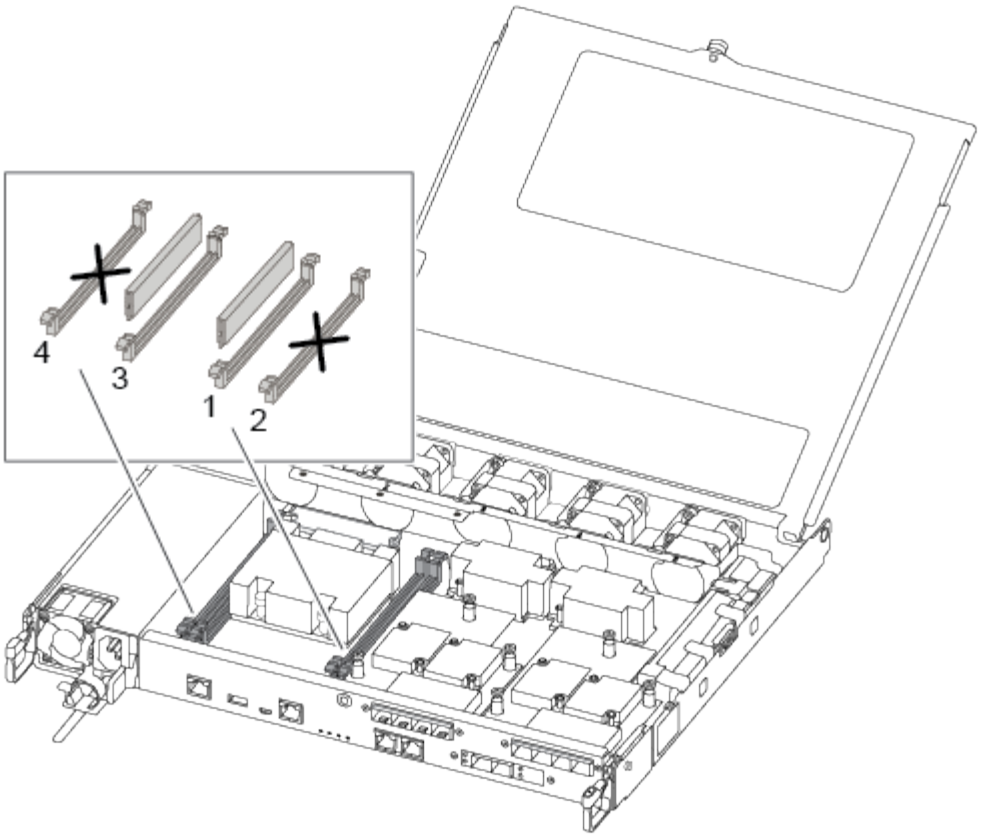

2. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
3. Slowly push apart the DIMM ejector tabs on either side of the DIMM, and slide the DIMM out of the slot.
4. Leave DIMM ejector tabs on the connector in the open position.
5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   Hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

6. Insert the replacement DIMM squarely into the slot.

   The DIMMs fit tightly in the socket. If not, reinsert the DIMM to realign it with the socket.

7. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis.

You can use the following illustration or the written steps to install the replacement controller module in the chassis.

1. If you have not already done so, install the air duct.

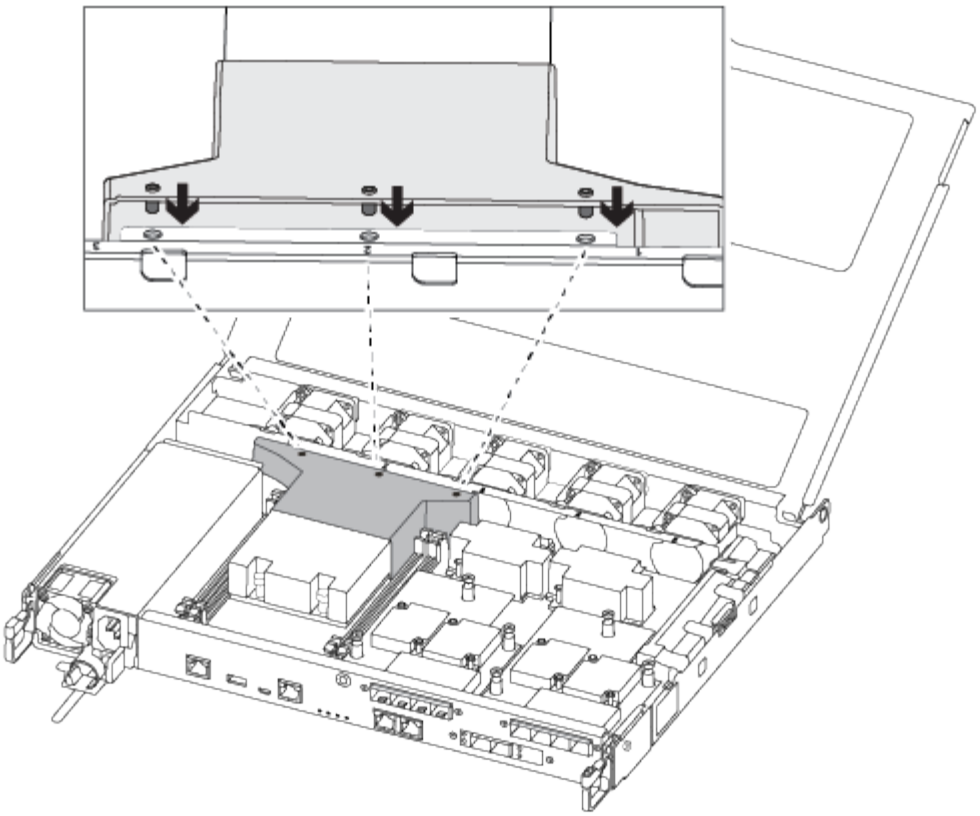

2. Close the controller module cover and tighten the thumbscrew.

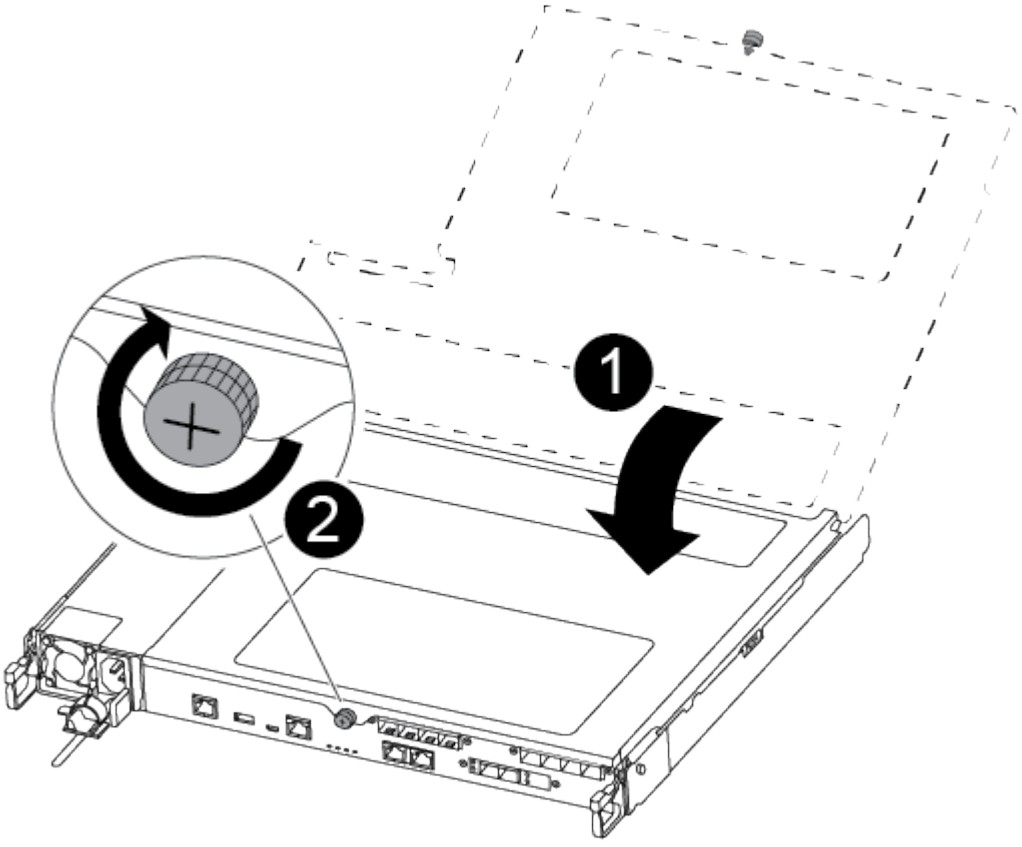

1	Controller module cover
2	Thumbscrew

3. Insert the controller module into the chassis:

   a. Ensure the latching mechanism arms are locked in the fully extended position.

   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.

   c. Place your index fingers through the finger holes from the inside of the latching mechanism.

   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.

   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   The controller module should be fully inserted and flush with the edges of the chassis.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace SSD Drive or HDD Drive - AFF C190

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

   When the drive’s activity LED is solid, it means that the drive has power. When the drive’s activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through 7.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

13. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan — FAS500f

You replace a fan with a new fan module when it fails.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a fan module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

   If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

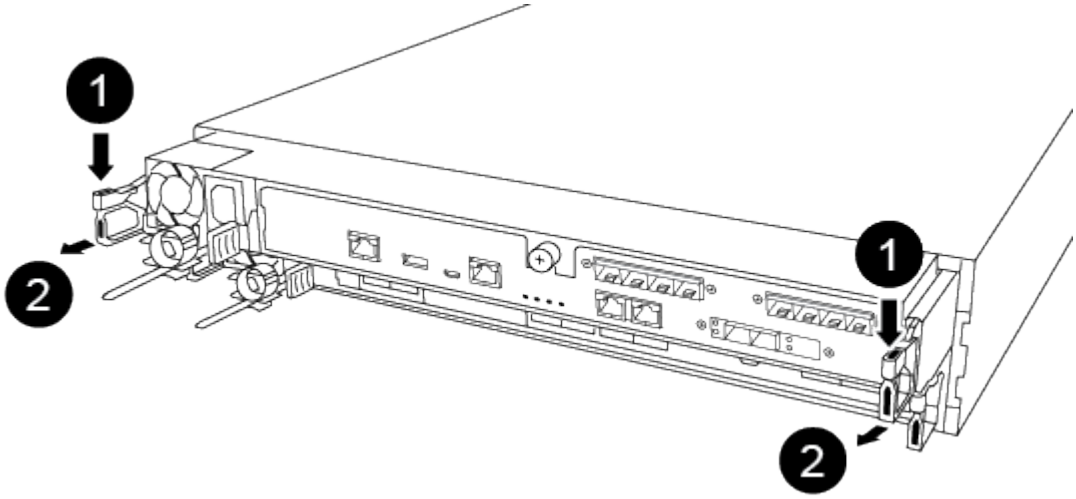

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat,

stable surface.

6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

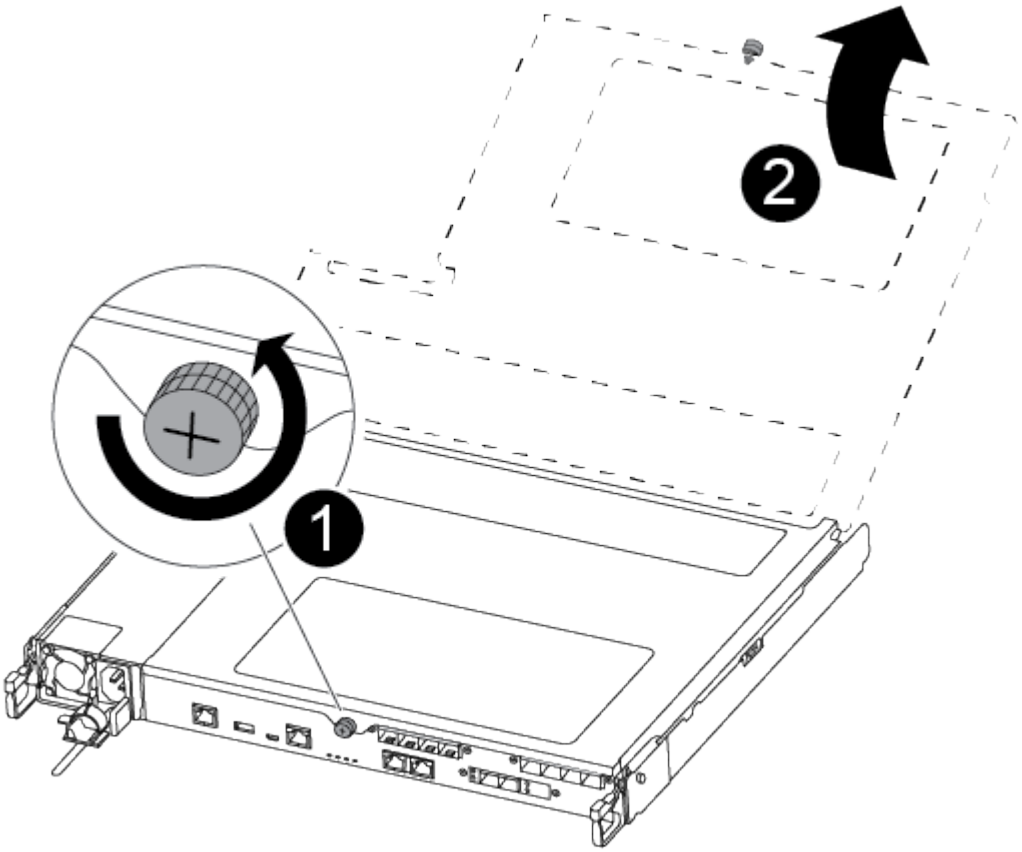

1	Thumbscrew
2	Controller module cover

Step 3: Replace a fan

To replace a fan, remove the failed fan module and replace it with a new fan module.

You can use the following video or the tabulated steps to replace a fan:

Animation - Replace a fan

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
2. Remove the fan module by pinching the side of the fan module, and then lifting the fan module straight out of the controller module.

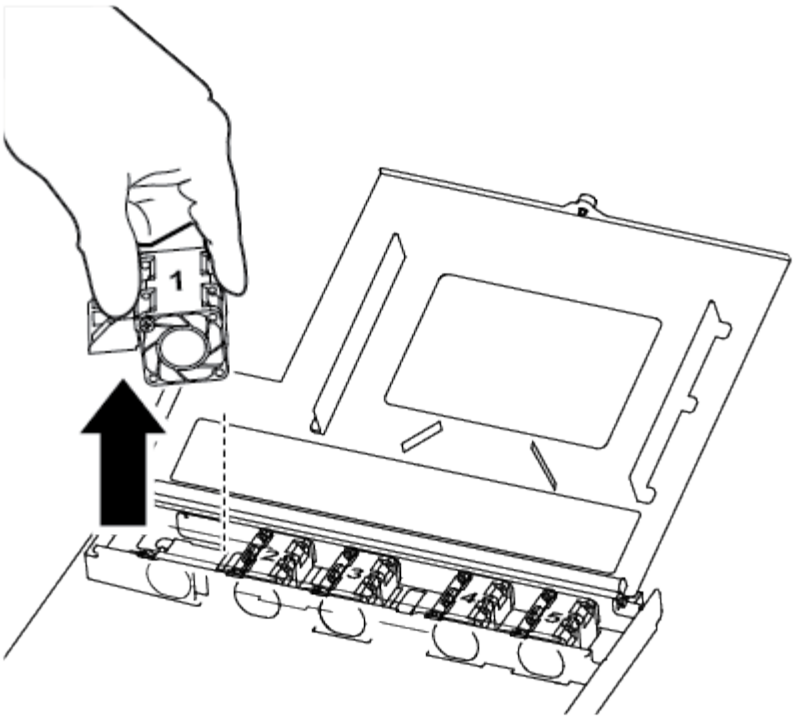

1	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Close the controller module cover and tighten the thumbscrew.

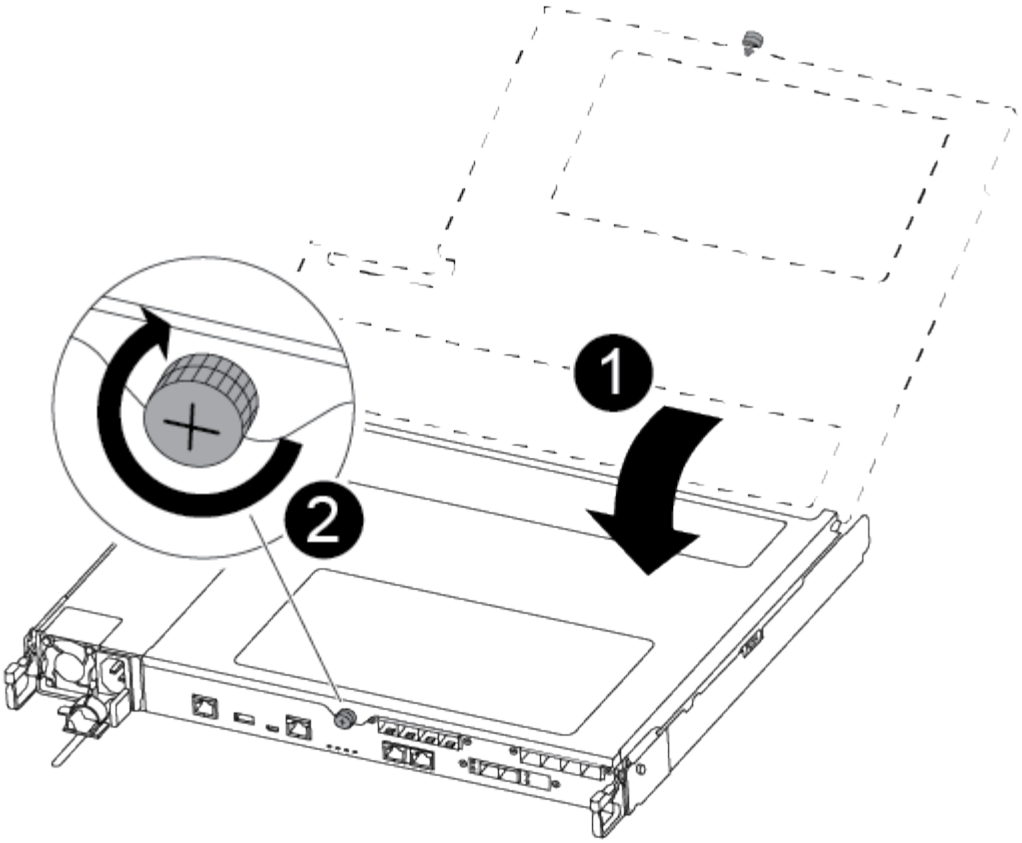

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:

   a. Ensure the latching mechanism arms are locked in the fully extended position.

   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.

   c. Place your index fingers through the finger holes from the inside of the latching mechanism.

   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.

   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   The controller module should be fully inserted and flush with the edges of the chassis.

3. Recable the system, as needed.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace or install a mezzanine card - FAS500f

To replace a failed mezzanine card, you must remove the cables and any SFP or QSFP modules, replace the card, reinstall the SFP or QSFP modules and recable the cards. To install a new mezzanine card, you must have the appropriate cables and SFP or QSFP modules.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

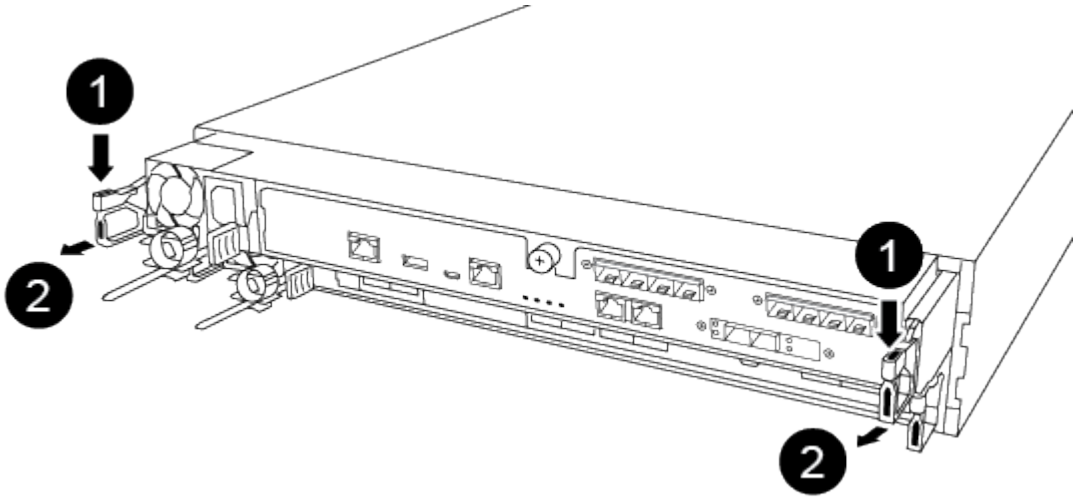

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

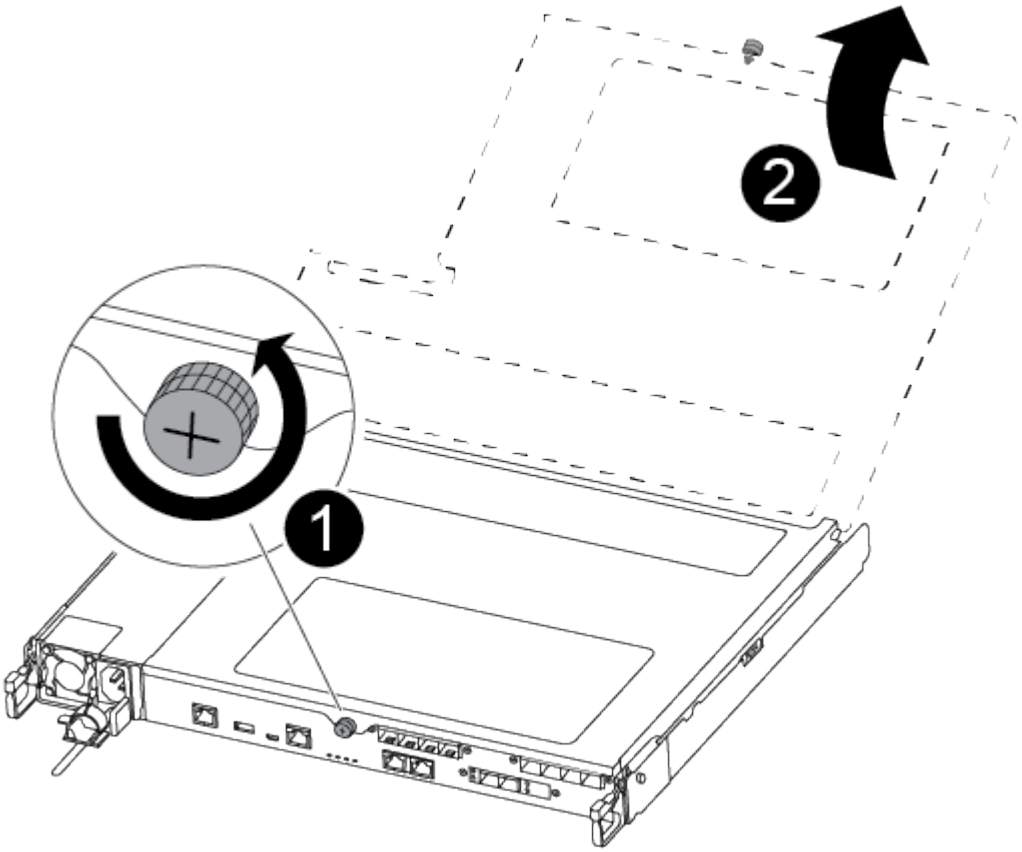

1	Thumbscrew
2	Controller module cover.

Step 3: Replace or install a mezzanine card

To replace a mezzanine card, you must remove the impaired card and install the replacement card; to install a mezzanine card, you must remove the faceplate and install the new card.

You can use the following video or the tabulated steps to replace a mezzanine card:

Animation - Replace a mezzanine card

Option 1: Replace a mezzanine card:

1. Locate and replace the impaired mezzanine card on your controller module.

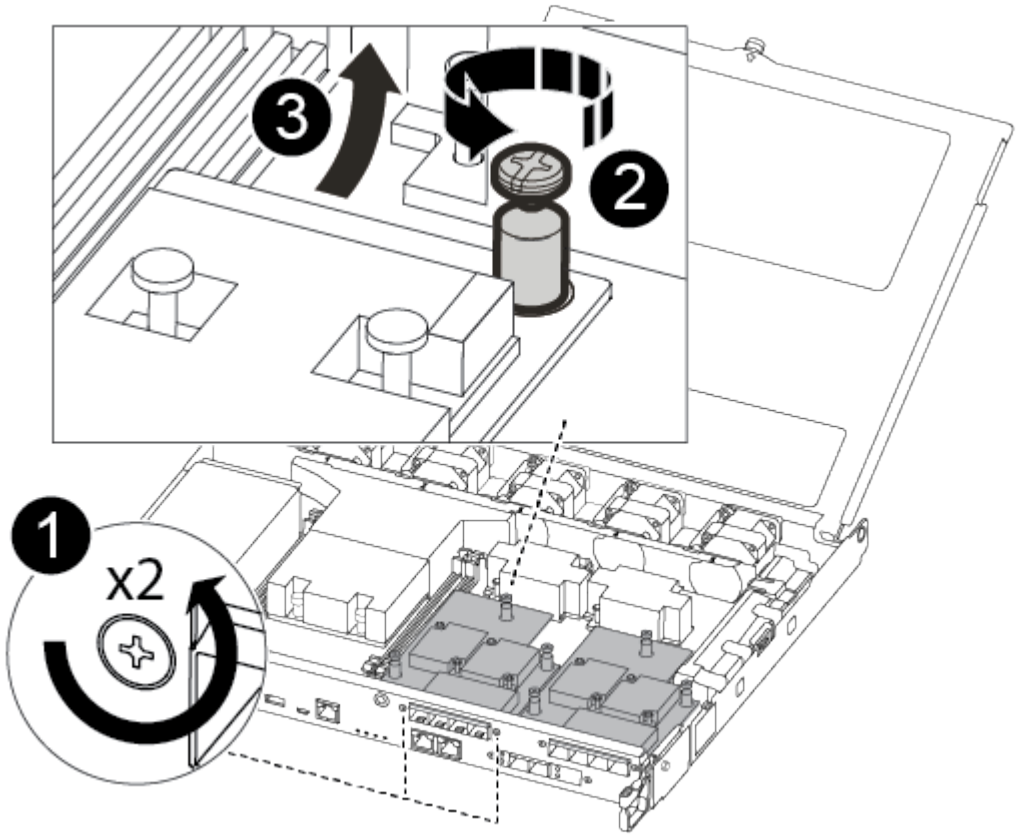

1	Remove screws on the face of the controller module.
2	Loosen the screw in the controller module.
3	Remove the mezzanine card.

2. Unplug any cabling associated with the impaired mezzanine card.

   Make sure that you label the cables so that you know where they came from.

3. Remove any SFP or QSFP modules that might be in the impaired mezzanine card and set it aside.
4. Using the #1 magnetic screwdriver, remove the screws from the face of the controller module and set them aside safely on the magnet.
5. Using the #1 magnetic screwdriver, loosen the screw on the impaired mezzanine card.
6. Using the #1 magnetic screwdriver, gently lift the impaired mezzanine card directly out of the socket and set it aside.
7. Remove the replacement mezzanine card from the antistatic shipping bag and align it to the inside face of the controller module.
8. Gently align the replacement mezzanine card into place.

9. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the controller module and on the mezzanine card.

   Do not apply force when tightening the screw on the mezzanine card; you might crack it.

10. Insert any SFP or QSFP modules that were removed from the impaired mezzanine card to the replacement mezzanine card.

Option 2: Install a mezzanine card:

You install a new mezzanine card if your system does not have one.

1. Using the #1 magnetic screwdriver, remove the screws from the face of the controller module and the faceplate covering the mezzanine card slot, and set them aside safely on the magnet.
2. Remove the mezzanine card from the antistatic shipping bag and align it to the inside face of the controller module.
3. Gently align the mezzanine card into place.
4. Using the #1 magnetic screwdriver, insert and tighten the screws on the face of the controller module and on the mezzanine card.

   Do not apply force when tightening the screw on the mezzanine card; you might crack it.

Step 4: Reinstall the controller module

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Close the controller module cover and tighten the thumbscrew.

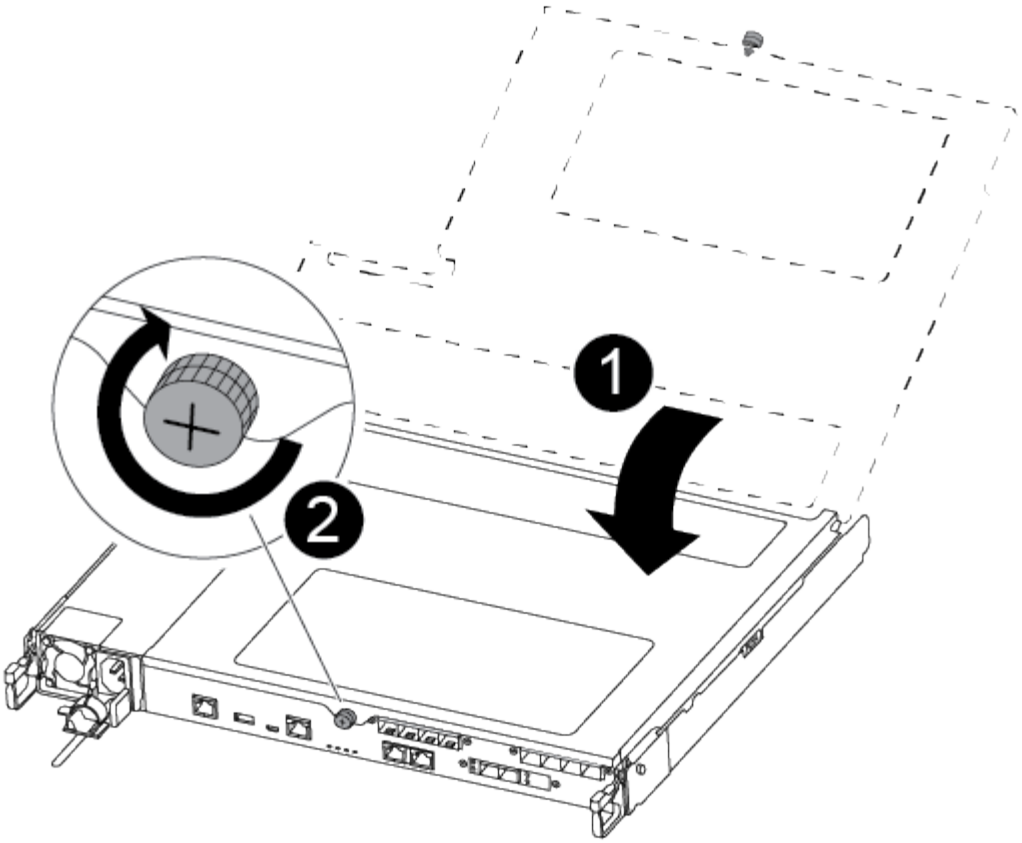

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis

   a. Ensure the latching mechanism arms are locked in the fully extended position.

   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.

   c. Place your index fingers through the finger holes from the inside of the latching mechanism.

   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.

   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   The controller module should be fully inserted and flush with the edges of the chassis.

3. Recable the system, as needed.

4. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVMEM battery - FAS500f

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

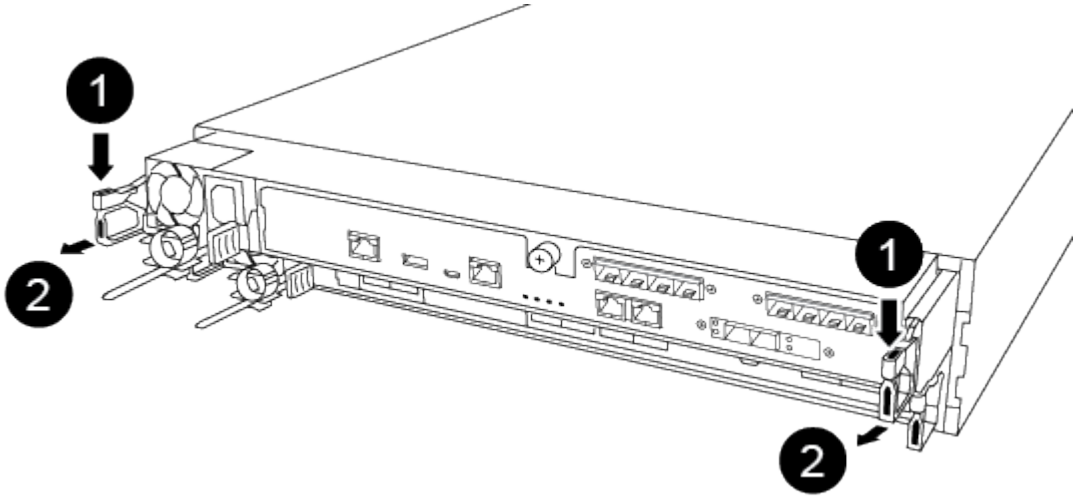

1	Lever

2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.
6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

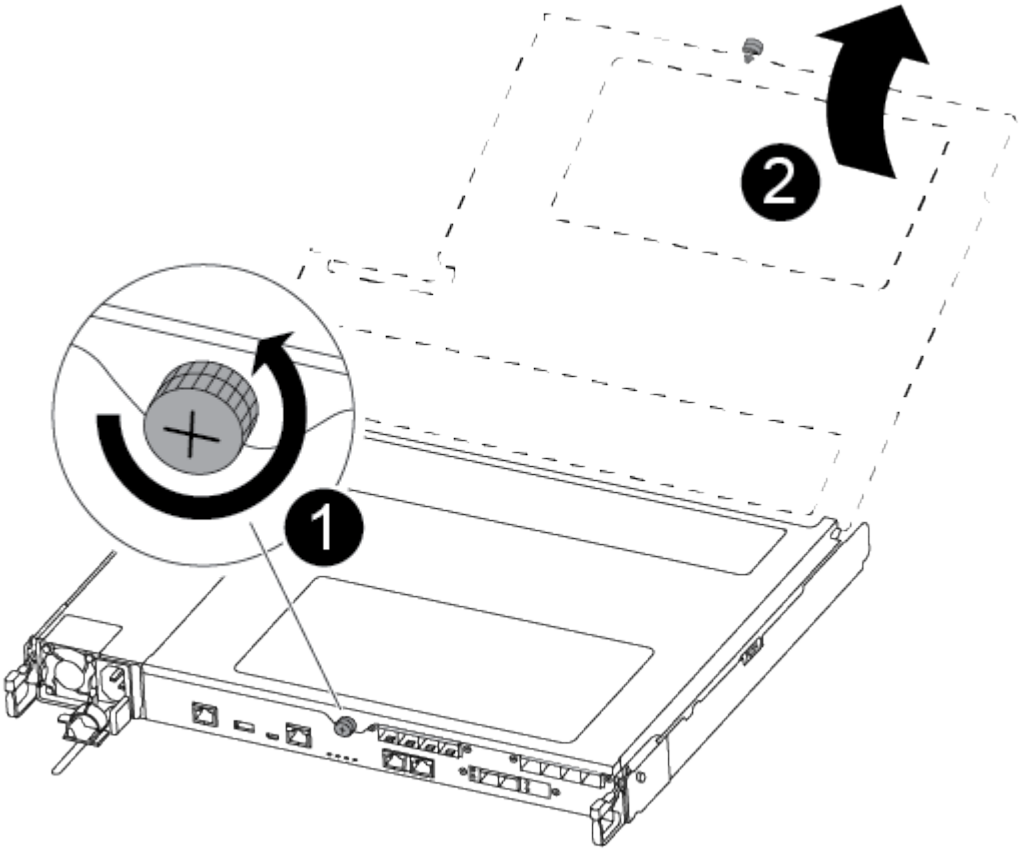

1	Thumbscrew
2	Controller module cover.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module.

You can use the following video or the tabulated steps to replace the NVMEM battery:

Animation - Replace the NVMEM battery

1. Locate and replace the impaired NVMEM battery on your controller module.

It is recommended that you follow the illustrated instructions in the order listed.

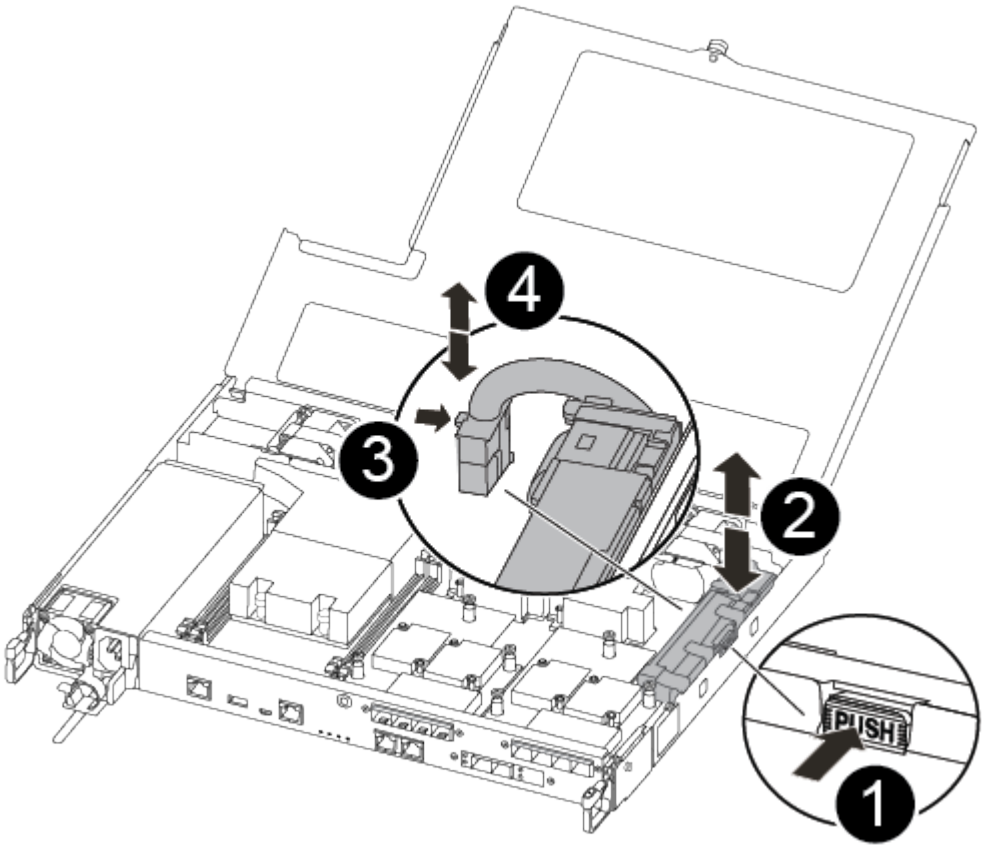

1	Squeeze the clip on the face of the battery plug.
2	Unplug the battery cable from the socket.
3	Grasp the battery and press the blue locking tab marked PUSH.
4	Lift the battery out of the holder and controller module.

2. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket.
3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module and set it aside.
4. Remove the replacement NV battery from the antistatic shipping bag and align it to the battery holder.
5. Insert the replacement NV battery plug into the socket.
6. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
7. Press firmly down on the battery pack to make sure that it is locked into place.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following illustration or the written steps to install the replacement controller module in the chassis.

1. Close the controller module cover and tighten the thumbscrew.

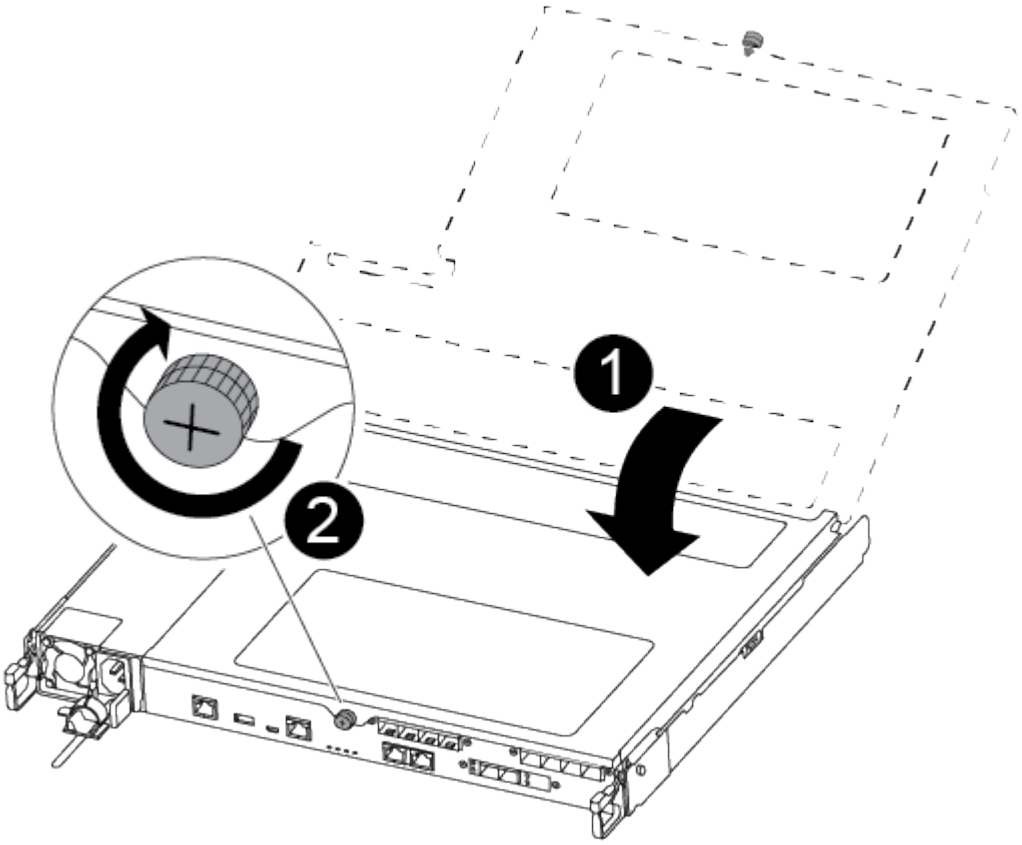

1	Controller module cover
2	Thumbscrew

2. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

The controller module begins to boot as soon as it is fully seated in the chassis.

The controller module should be fully inserted and flush with the edges of the chassis.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

> You will connect the rest of the cables to the controller module later in this procedure.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a power supply - FAS500f

Replacing a power supply involves disconnecting the target power supply (PSU) from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting it to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

You can use the following video or the tabulated steps to replace the power supply:

Animation - Replace the power supply

1. If you are not already grounded, properly ground yourself.
2. Identify the power supply you want to replace, based on console error messages or through the red Fault LED on the power supply.
3. Disconnect the power supply:
   a. Open the power cable retainer, and then unplug the power cable from the power supply.
   b. Unplug the power cable from the power source.
4. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

> The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

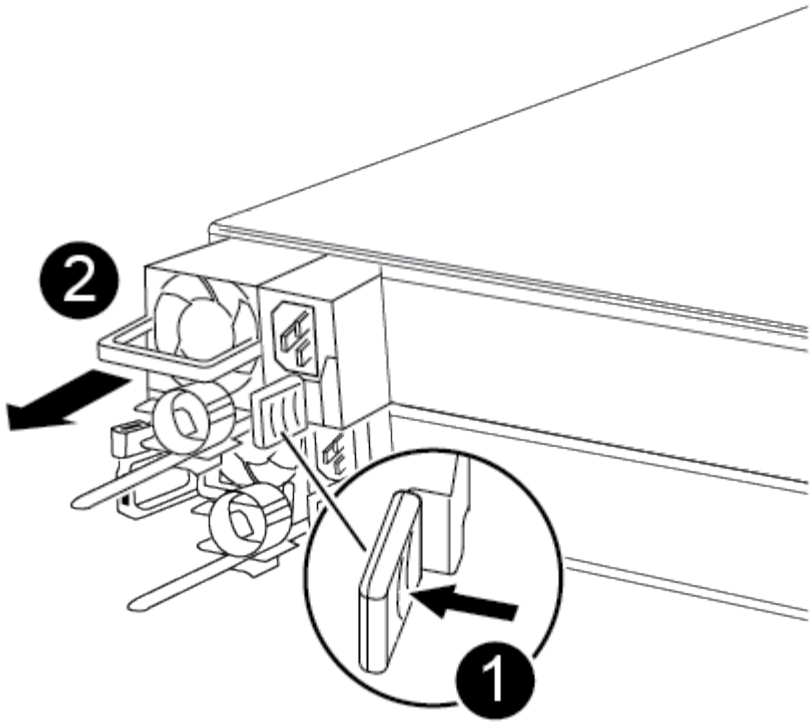

1	Blue power supply locking tab
2	Power supply

5. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

6. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply and the power source.

   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization

continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace a component inside the controller module.

Make sure that you label the cables so that you know where they came from.

1. If you are not already grounded, properly ground yourself.
2. Unplug the controller module power supplies from the source.
3. Release the power cable retainers, and then unplug the cables from the power supplies.
4. Insert your forefinger into the latching mechanism on either side of the controller module, press the lever with your thumb, and gently pull the controller a few inches out of the chassis.

> If you have difficulty removing the controller module, place your index fingers through the finger holes from the inside (by crossing your arms).

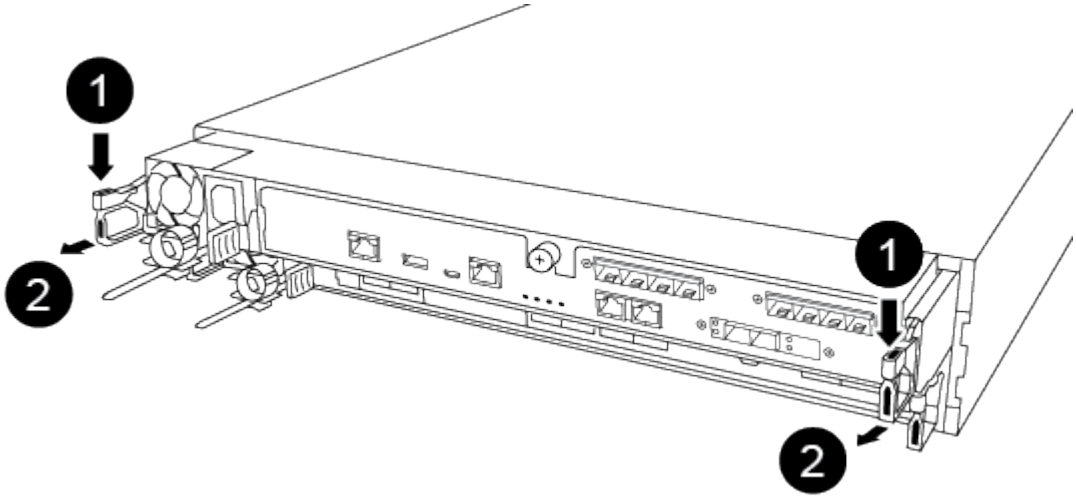

1	Lever
2	Latching mechanism

5. Using both hands, grasp the controller module sides and gently pull it out of the chassis and set it on a flat, stable surface.

6. Turn the thumbscrew on the front of the controller module anti-clockwise and open the controller module cover.

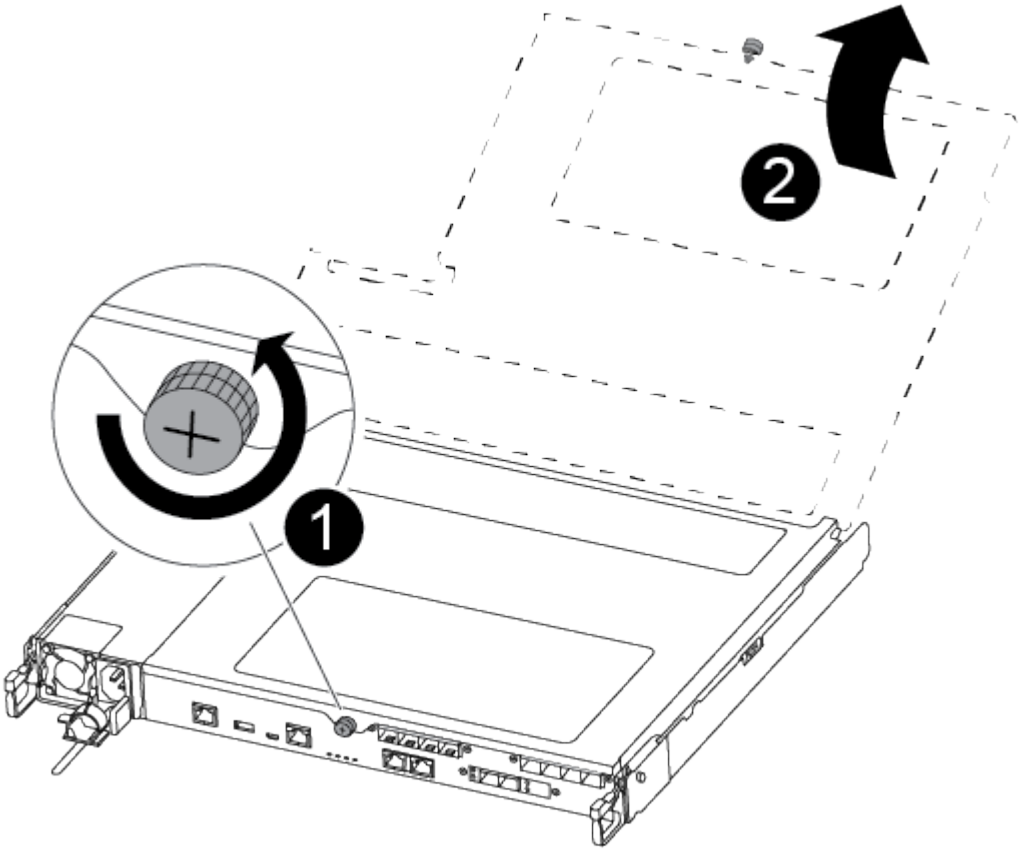

1	Thumbscrew
2	Controller module cover.

7. Lift out the air duct cover.

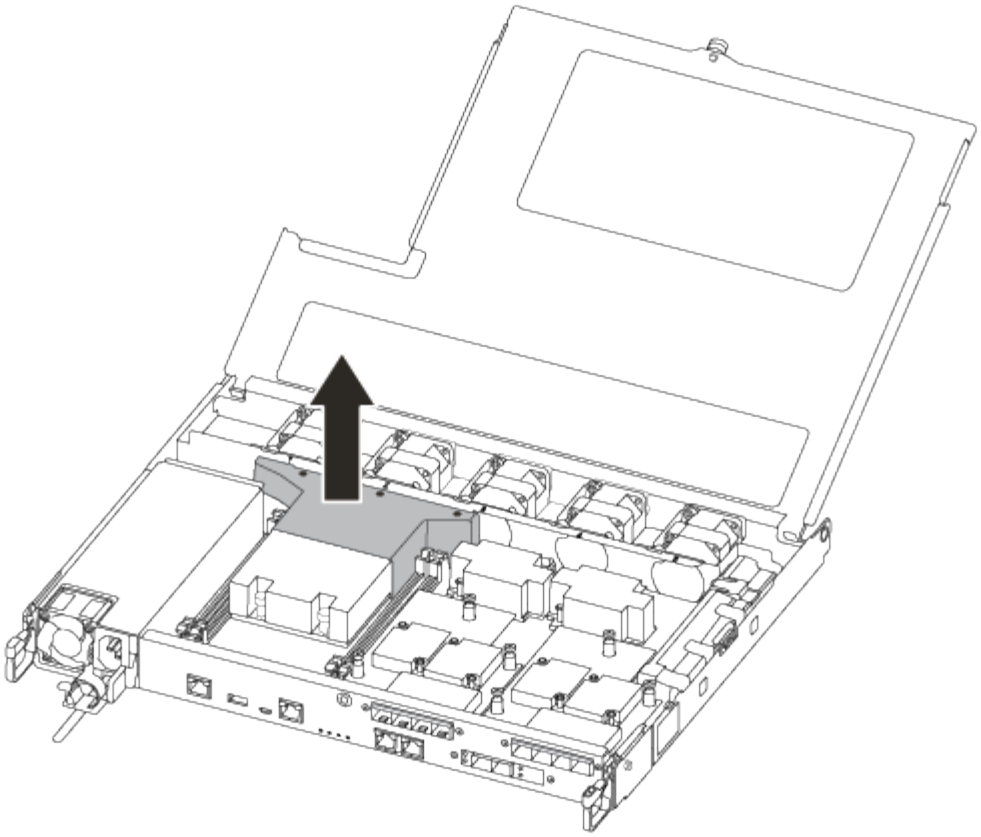

Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

You can use the following video or the tabulated steps to replace the RTC battery:

Animation - Replace the RTC battery

1. Locate the RTC battery between the heatsink and the midplane and remove it exactly as shown in the graphic.

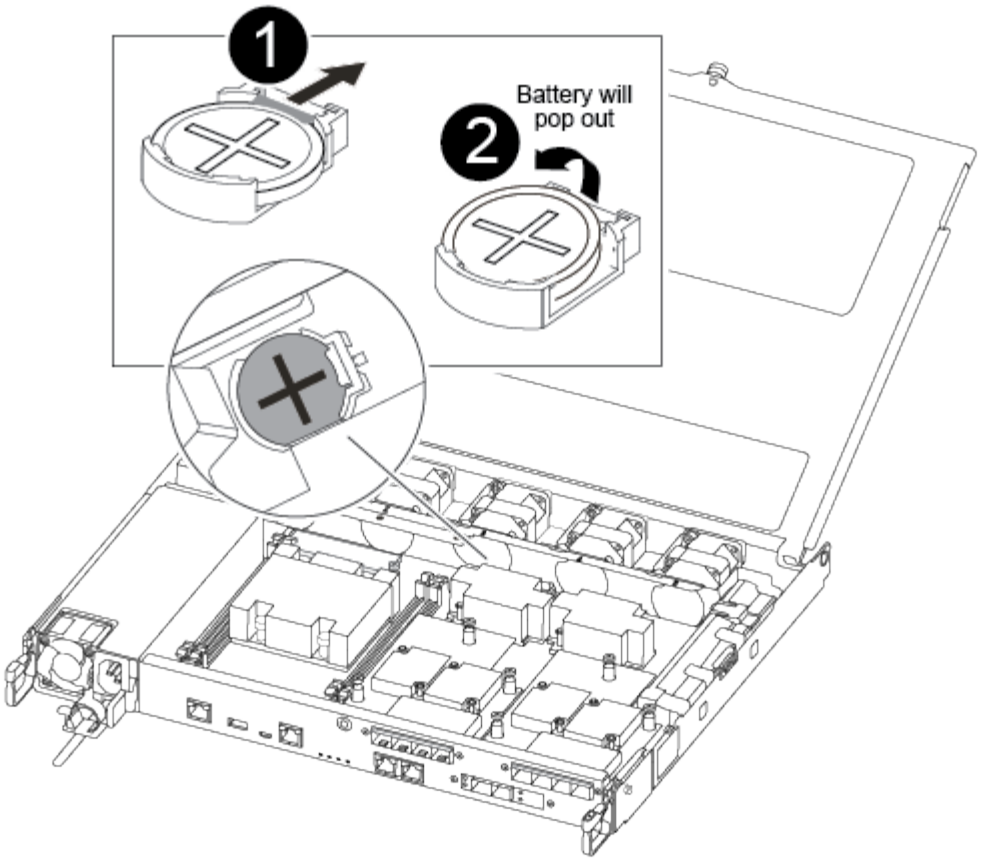

1	Gently pull tab away from the battery housing. NOTE: Pulling it away aggressively might displace the tab.
2	Lift the battery up. Make a note of the polarity of the battery.
3	The battery should eject out.

The battery will be ejected out.

2. Remove the replacement battery from the antistatic shipping bag.
3. Locate the RTC battery holder between the heatsink and the midplane and insert it exactly as shown in the graphic.

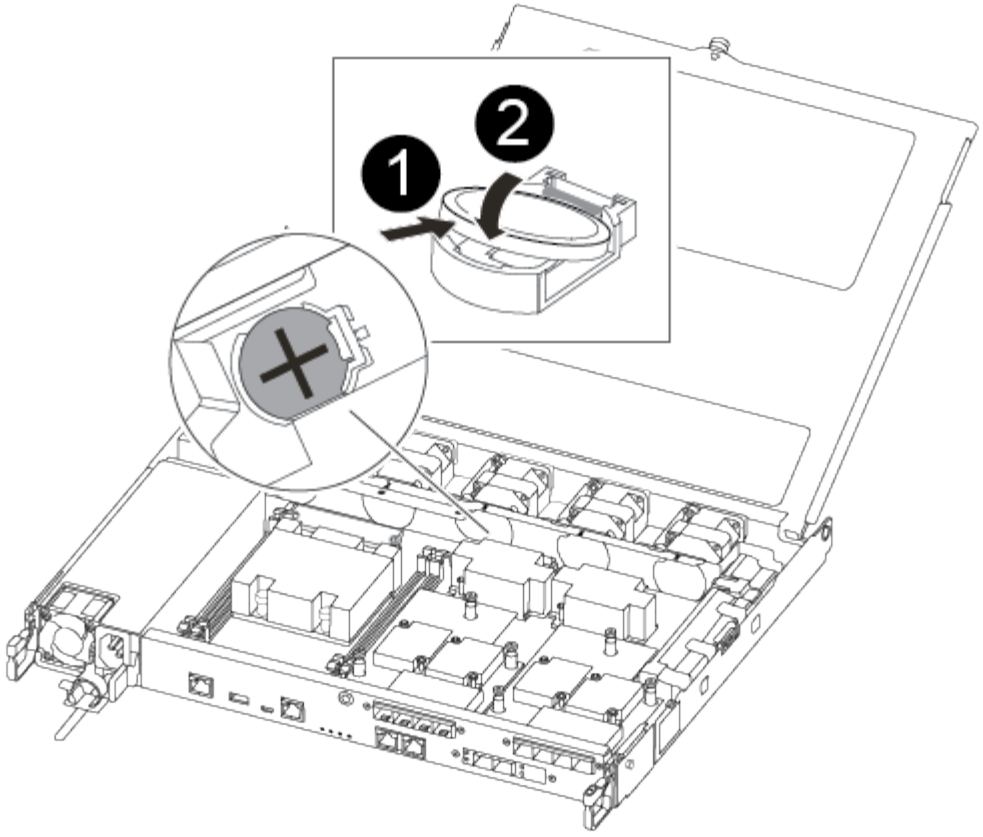

1	With positive polarity face up, slide the battery under the tab of the battery housing.
2	Push the battery gently into place and make sure the tab secures it to the housing. Pushing it in aggressively might cause the battery to eject out again.

4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and set time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

Steps

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Insert the controller module into the chassis:
   a. Ensure the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, align and gently slide the controller module into the latching mechanism arms until it stops.
   c. Place your index fingers through the finger holes from the inside of the latching mechanism.
   d. Press your thumbs down on the orange tabs on top of the latching mechanism and gently push the controller module over the stop.
   e. Release your thumbs from the top of the latching mechanisms and continue pushing until the latching mechanisms snap into place.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   f. Halt the controller at the LOADER prompt.

   The controller module should be fully inserted and flush with the edges of the chassis.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Complete the replacement process

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

FAS8200 systems

Install and setup

Cluster configuration worksheet - FAS8200

You can use the Cluster Configuration Worksheet to gather and record your site-specific

IP addresses and other information required when configuring an ONTAP cluster.

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

For MetroCluster configurations, see either:

- Install MetroCluster IP configuration
- Install MetroCluster Fabric-Attached configuration

Installation and setup PDF poster - FAS8200

You can use the PDF poster to install and set up your new system. The AFF FAS8200 Installation and Setup Instructions provides step-by-step instructions with live links to additional content.

Maintain

Maintain FAS8200 hardware

For the FAS8200 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Caching module

You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

NVEM battery

A battery is included with a controller and preserves cached data if the AC power fails.

PCIe card

A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - FAS8200

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct node:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy node* is the HA partner of the impaired node.

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the impaired controller - FAS8200

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.
After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster configuration

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.

If the impaired controller is displaying…	Then…
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 3: Controller is in a two-node Metrocluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes

that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

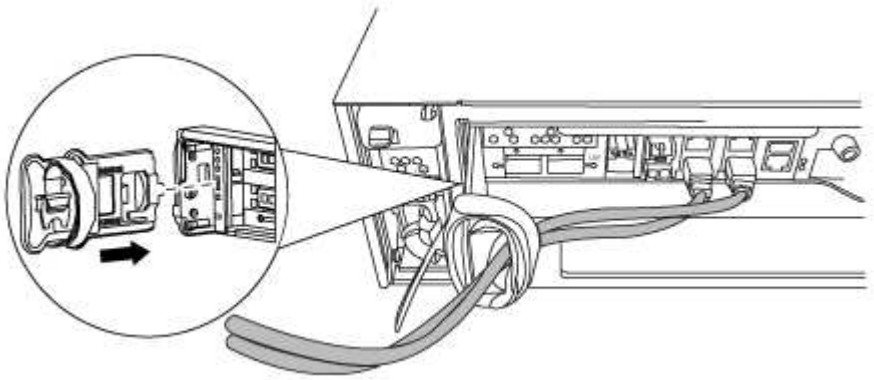

4. Loosen the thumbscrew on the cam handle on the controller module.

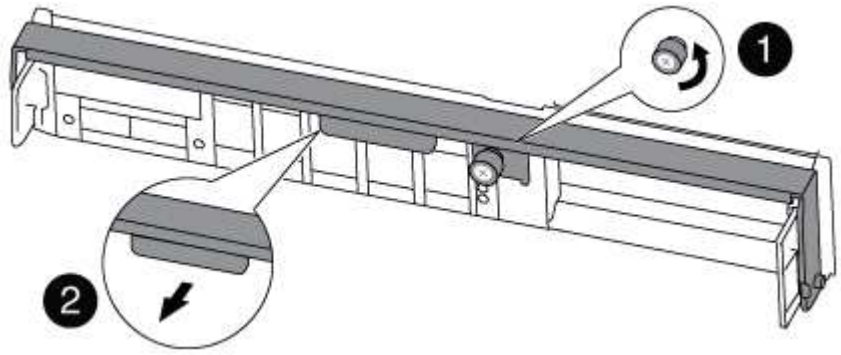

1	
Thumbscrew	
2	

Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

1. If you are not already grounded, properly ground yourself.
2. Locate the boot media using the following illustration or the FRU map on the controller module:

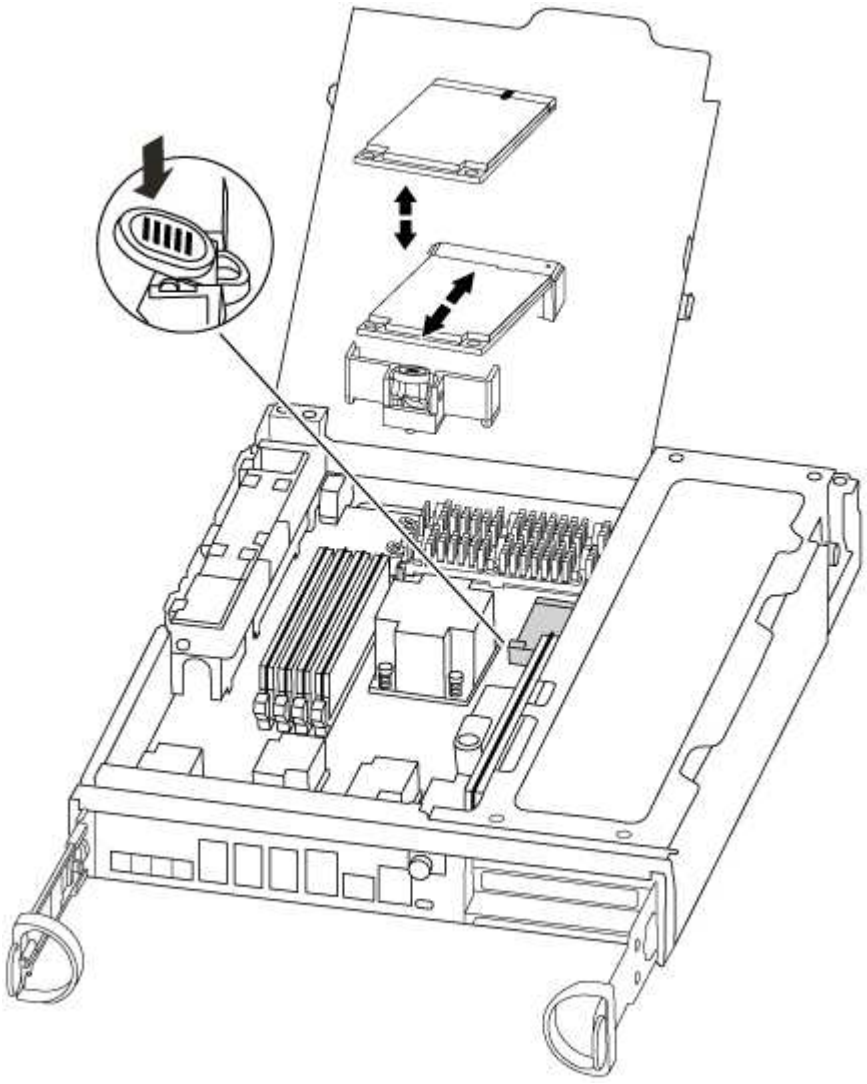

3. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

4. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
5. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

6. Push the boot media down to engage the locking button on the boot media housing.
7. Close the controller module cover.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.

  The tar.gz file must be placed on a FAT32 formatted partition that is a minimum of 4GB. While FAT32 partitions can be as large as 2TB, Windows built-in tools (e.g. DIskPart) cannot format FAT32 partitions larger than 32GB.

- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs) if they were removed.

3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

   The controller begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt

the controller to boot to LOADER.

6. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

   The system begins to boot and stops at the LOADER prompt.

7. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

     The target port you configure is the target port you use to communicate with the impaired controller from the healthy controller during var file system restore with a network connection. You can also use the e0M port in this command.

   - If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
     - filer_addr is the IP address of the storage system.
     - netmask is the network mask of the management network that is connected to the HA partner.
     - gateway is the gateway for the network.
     - dns_addr is the IP address of a name server on your network.
     - dns_domain is the Domain Name System (DNS) domain name.

       If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

     Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

8. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

The changes will be implemented when the system is booted.

Boot the recovery image - FAS8200

The procedure for booting the impaired controller from the recovery image depends on whether the system is in a two-node MetroCluster configuration.

Option 1: Most systems

:
You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

This procedure applies to systems that are not in a two-node MetroCluster configuration.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:
   - If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed
   - If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.
6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.
8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Option 2: Controller is in a two-node MetroCluster

You must boot the ONTAP image from the USB drive and verify the environmental variables.

This procedure applies to systems in a two-node MetroCluster configuration.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. After the image is installed, start the restoration process:
   a. Press `n` when prompted to restore the backup configuration.
   b. Press `y` when prompted to reboot to start using the newly installed software.

      You should be prepared to interrupt the boot process when prompted.

4. As the system boots, press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message., and when the Boot Menu is displayed select option 6.
5. Verify that the environmental variables are set as expected.
   a. Take the node to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv *environment-variable-name changed-value*` command.

d. Save your changes using the `savenv` command.

e. Reboot the node.

Switch back aggregates in a two-node MetroCluster configuration - FAS8200

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    ------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Restore encryption - FAS8200

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** ``` Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10 ```

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   `storage failover modify -node local -auto-giveback true`

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   `system node autosupport invoke -node * -type all -message MAINT=END`

Return the failed part to NetApp - FAS8200

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the caching module - FAS8200

You must replace the caching module in the controller module when your system registers a single AutoSupport (ASUP) message that the module has gone offline; failure to do so results in performance degradation.

You might want to erase the contents of your caching module before replacing it.

- Although data on the caching module is encrypted, you might want to erase any data from the impaired caching module and verify that the caching module has no data:
  - Erase the data on the caching module: `system controller flash-cache secure-erase run`

`-node node name localhost -device-id device_number`

> Run the `system controller flash-cache show` command if you don’t know the Flash Cache device ID.

- Verify that the data has been erased from the caching module: `system controller flash-cache secure-erase show`

- You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Open the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

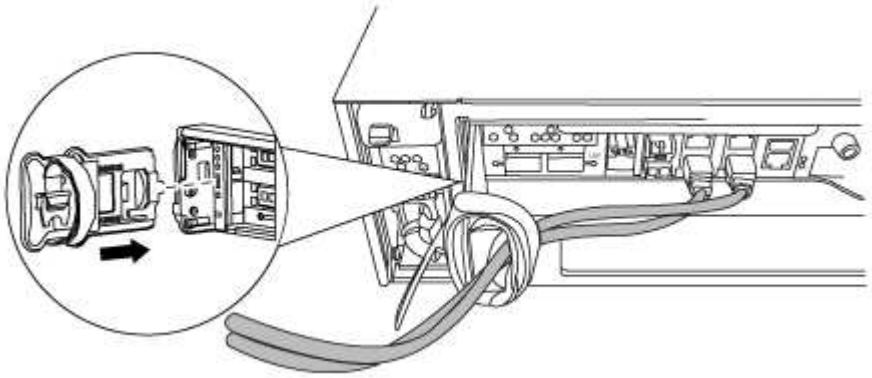

4. Loosen the thumbscrew on the cam handle on the controller module.

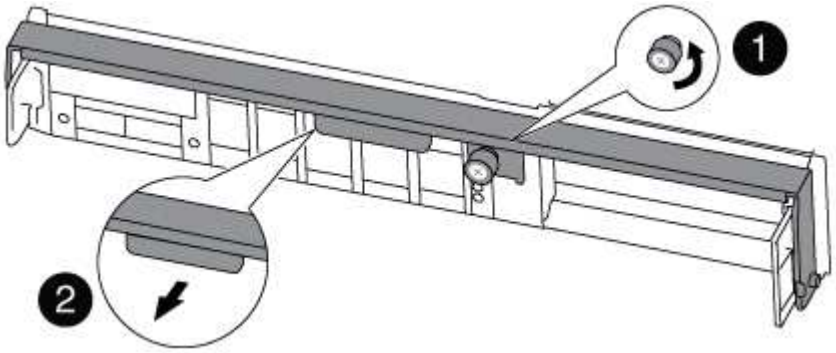

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace or add a caching module

To replace or add a caching module referred to as the M.2 PCIe card on the label on your controller, locate the slots inside the controller and follow the specific sequence of steps.

Your storage system must meet certain criteria depending on your situation:

- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- All other components in the storage system must be functioning properly; if not, you must contact technical support.

1. Locate the caching module at the rear of the controller module and remove it.
   a. Press the release tab.
   b. Remove the heatsink.

      The storage system comes with two slots available for the caching module and only one slot is occupied, by default.

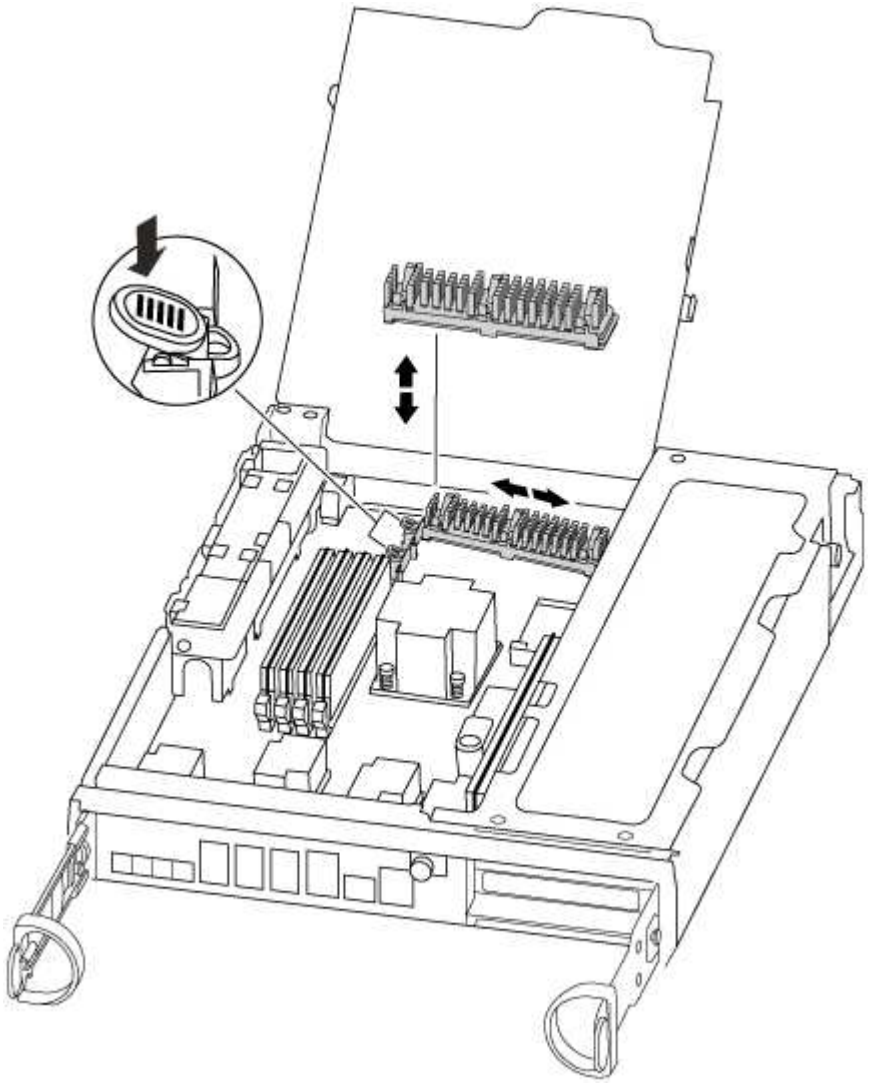

2. If you are adding a caching module, go to the next step; if you are replacing the caching module, gently pull it straight out of the housing.

3. Align the edges of the caching module with the socket in the housing, and then gently push it into the socket.
4. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

5. Reseat and push the heatsink down to engage the locking button on the caching module housing.
6. Repeat the steps if you have a second caching module. Close the controller module cover, as needed.

Step 4: Reinstall the controller

After you replace a component within the controller module, you must reinstall the controller module in the system chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   ⓘ Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

3. Complete the reinstallation of the controller module:

   The controller module begins to boot as soon as it is fully seated in the chassis

   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      ⓘ Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Tighten the thumbscrew on the cam handle on back of the controller module.
   c. If you have not already done so, reinstall the cable management device.
   d. Bind the cables to the cable management device with the hook and loop strap.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Complete the replacement process

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - FAS8200

To replace the chassis, you must move the power supplies, fans, and controller modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving the controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-controller cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - FAS8200

To replace the chassis, you must shutdown the controllers.

Option 1: Most configurations

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.

4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.

If the impaired controller…	Then…
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace hardware - FAS8200

Move the power supplies, fans, and controller modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Move a power supply

Moving out a power supply when replacing a chassis involves turning off, disconnecting, and removing the power supply from the old chassis and installing and connecting it on the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Press down the release latch on the power supply cam handle, and then lower the cam handle to the fully open position to release the power supply from the mid plane.

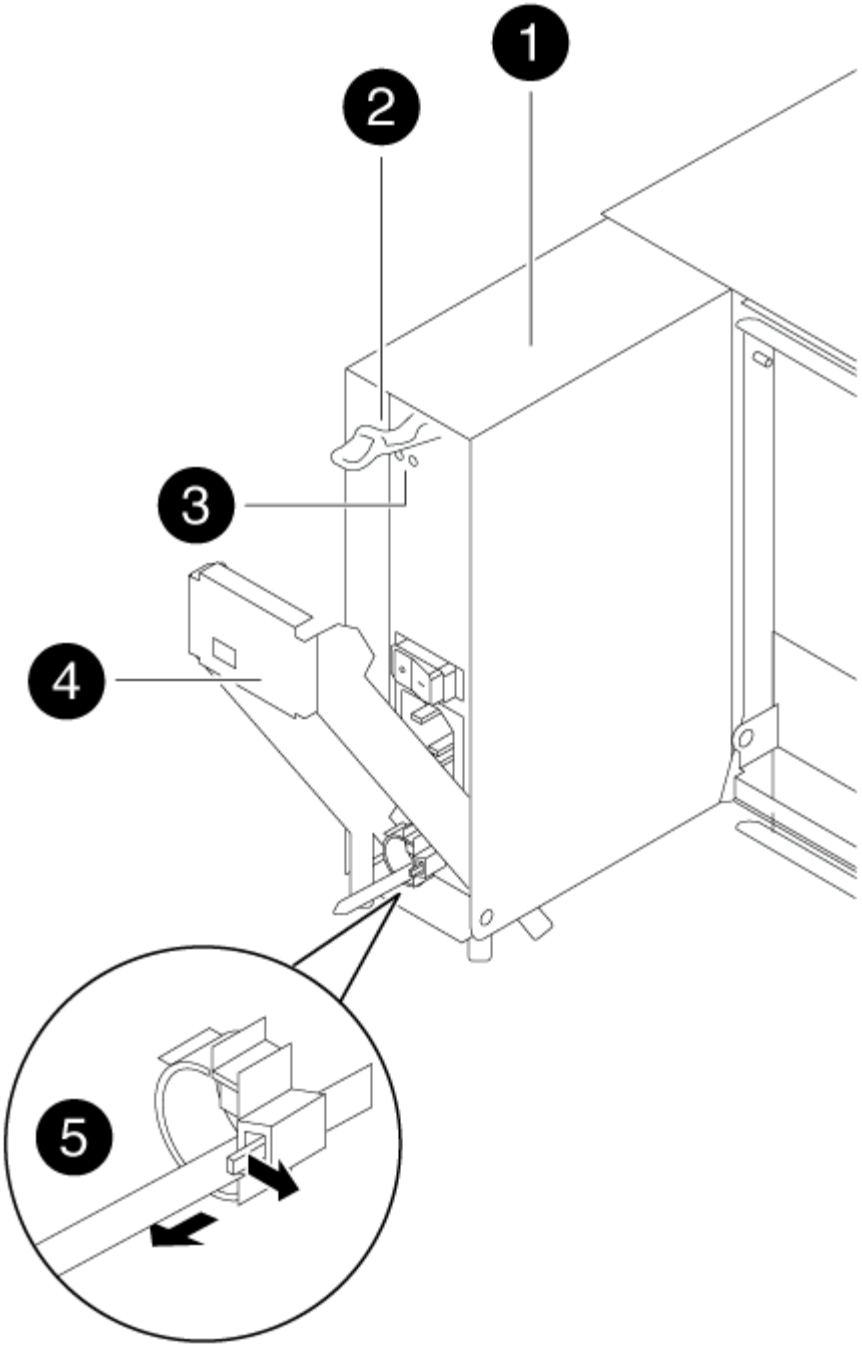

1	Power supply
2	Cam handle release latch
3	Power and Fault LEDs
4	Cam handle

5	Power cable locking mechanism

4. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Push firmly on the power supply cam handle to seat it all the way into the chassis, and then push the cam handle to the closed position, making sure that the cam handle release latch clicks into its locked position.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

   Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

Step 2: Move a fan

Moving out a fan module when replacing the chassis involves a specific sequence of tasks.

1. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
2. Press down the release latch on the fan module cam handle, and then pull the cam handle downward.

   The fan module moves a little bit away from the chassis.

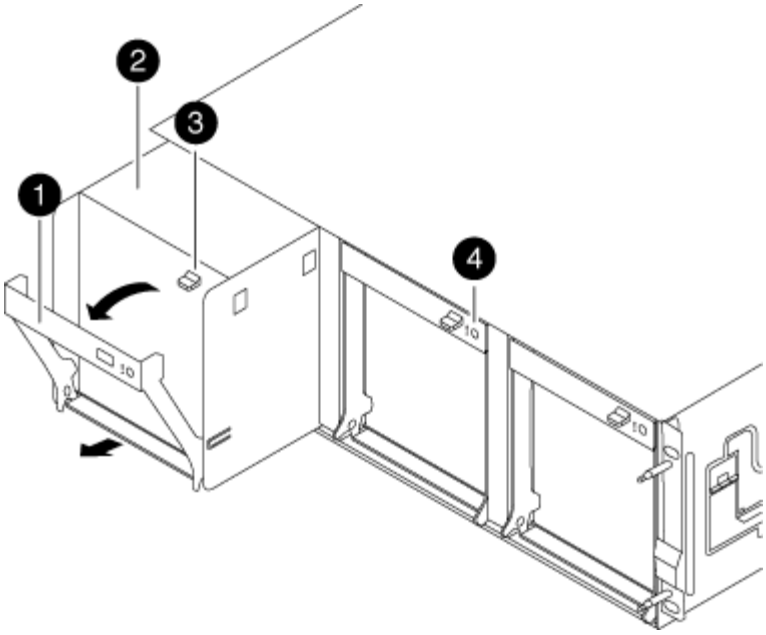

1	Cam handle
2	Fan module
3	Cam handle release latch
4	Fan module Attention LED

3. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

4. Set the fan module aside.
5. Repeat the preceding steps for any remaining fan modules.
6. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
7. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

8. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The fan LED should be green after the fan is seated and has spun up to operational speed.

9. Repeat these steps for the remaining fan modules.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Step 3: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the old chassis.

1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

2. Remove and set aside the cable management devices from the left and right sides of the controller module.

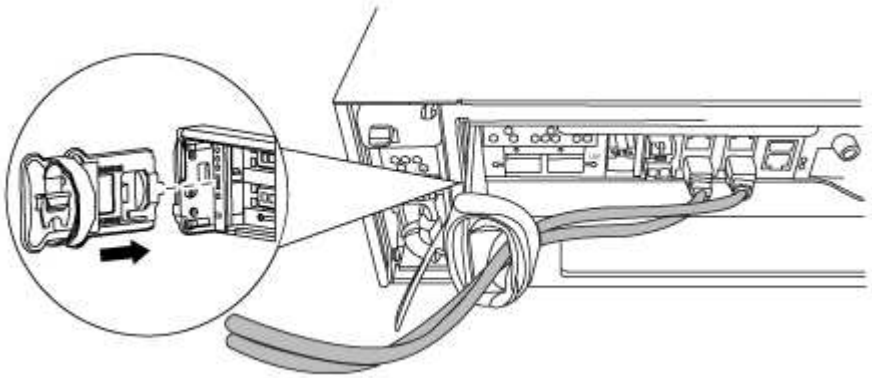

3. Loosen the thumbscrew on the cam handle on the controller module.

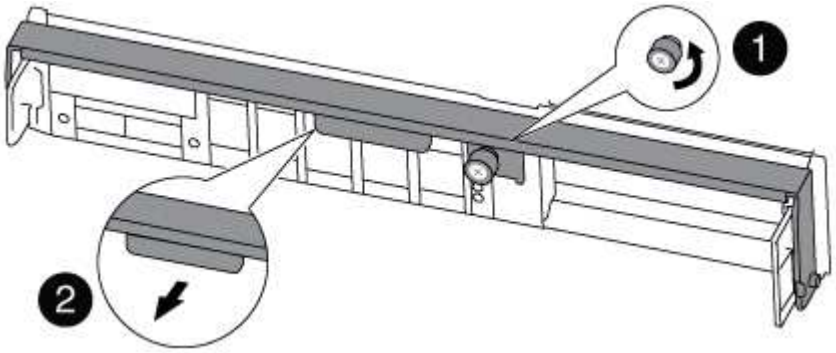

1	Thumbscrew
2	Cam handle

4. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 4: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.

   If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller

After you install the controller module and any other components into the new chassis, boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Repeat the preceding steps if there is a second controller to install in the new chassis.
4. Complete the installation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Repeat the preceding steps for the second controller module in the new chassis.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reinstall the blanking panel and then go to the next step.

5. Connect the power supplies to different power sources, and then turn them on.
6. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

      If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

   b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - FAS8200

You must verify the HA state of the chassis, switch back aggregates, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for HA-state can be one of the following:
      * **`ha`**
      * `mcc`
      * **`mcc-2n`**
      * `mccip`
      * `non-ha`

   b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

4. The next step depends on your system configuration.

If your system is in…	Then…
A stand-alone configuration	a. Exit Maintenance mode: `halt` b. Go to Completing the replacement process.
An HA pair with a second controller module	Exit Maintenance mode: `halt` The LOADER prompt appears.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - FAS8200

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system has a V_StorageAttach license, you must refer to the additional required steps before performing this procedure.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight controller MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- This procedure includes steps for automatically or manually reassigning drives to the *replacement* controller, depending on your system’s configuration.

  You should perform the drive reassignment as directed in the procedure.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- Any PCIe cards moved from the old controller module to the new controller module or added from existing customer site inventory must be supported by the replacement controller module.

  NetApp Hardware Universe

- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement* controller is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - FAS8200

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the controller module hardware - FAS8200

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Open the controller module

To replace the controller module, you must first remove the old controller module from the chassis.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

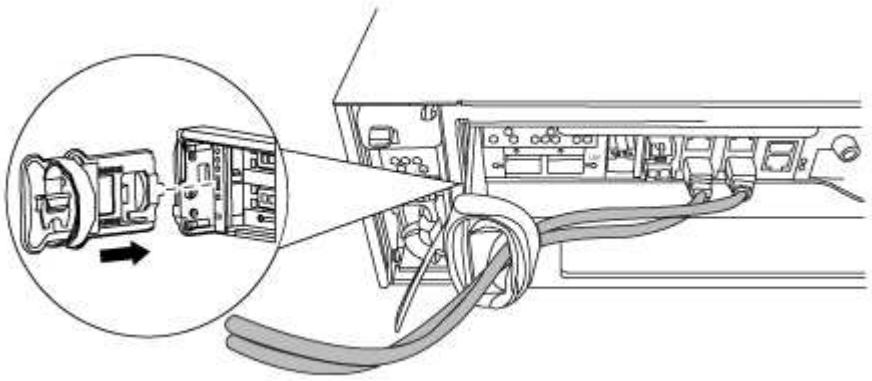

4. If you left the SFP modules in the system after removing the cables, move them to the new controller module.
5. Loosen the thumbscrew on the cam handle on the controller module.

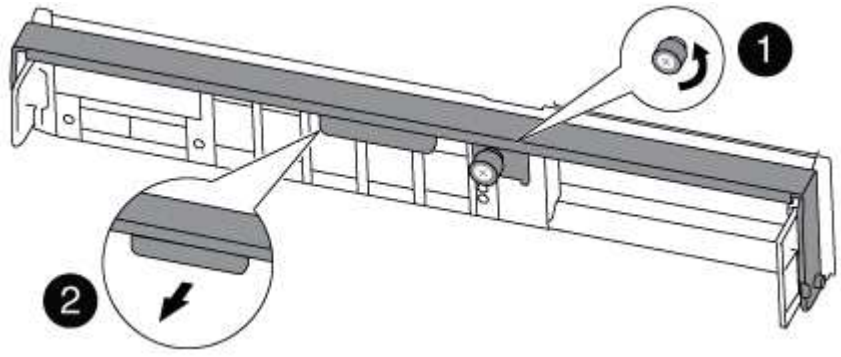

1	Thumbscrew
2	Cam handle

6. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 2: Move the boot device

You must locate the boot media and follow the directions to remove it from the old controller and insert it in the new controller.

1. Locate the boot media using the following illustration or the FRU map on the controller module:

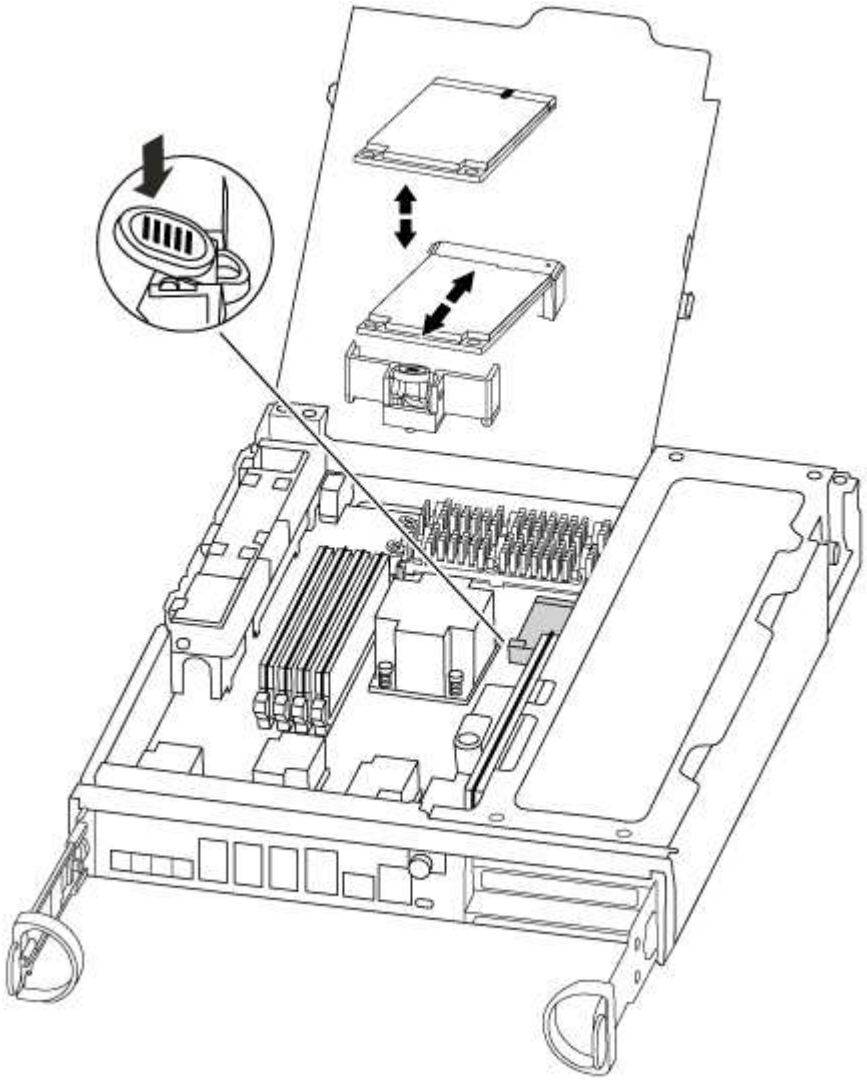

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the NVMEM battery

To move the NVMEM battery from the old controller module to the new controller module, you must perform a specific sequence of steps.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

     - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
     - If the LED is on and power is on, unwritten data is stored on NVMEM.

       This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Open the CPU air duct and locate the NVMEM battery.

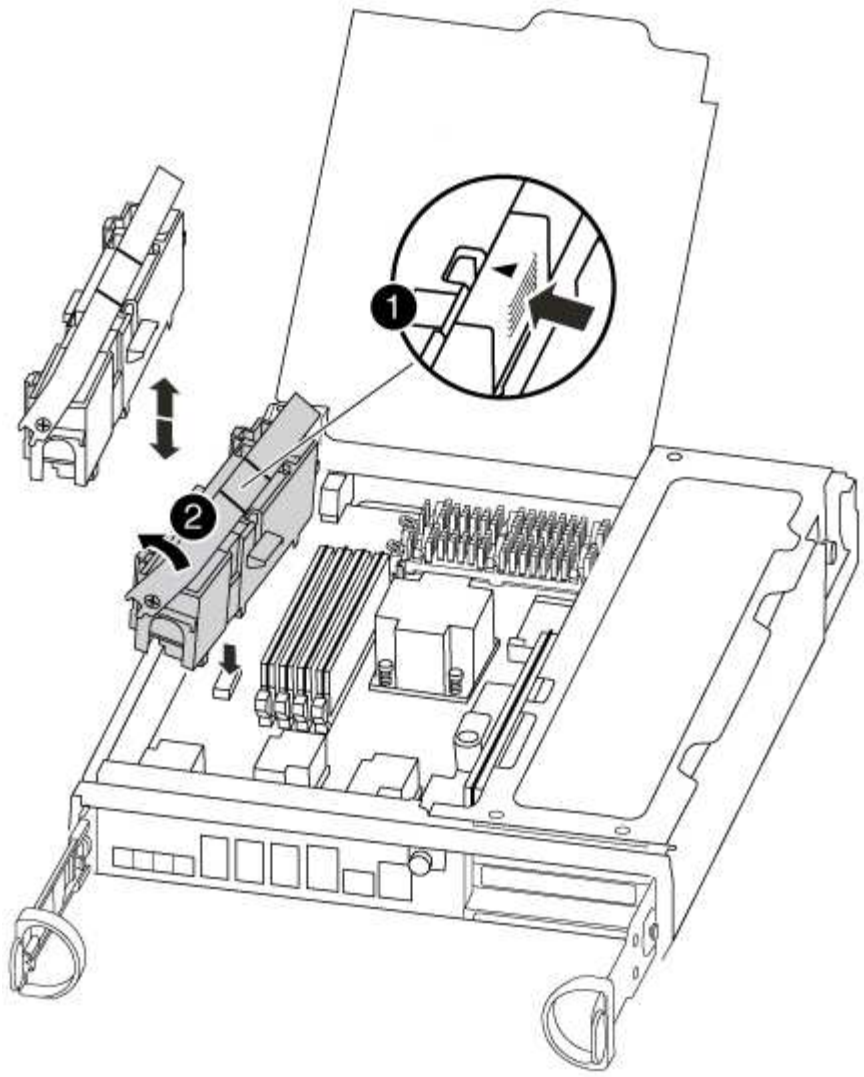

1	Battery lock tab
2	NVMEM battery pack

3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
4. Remove the battery from the controller module and set it aside.

Step 4: Move the DIMMs

To move the DIMMs, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

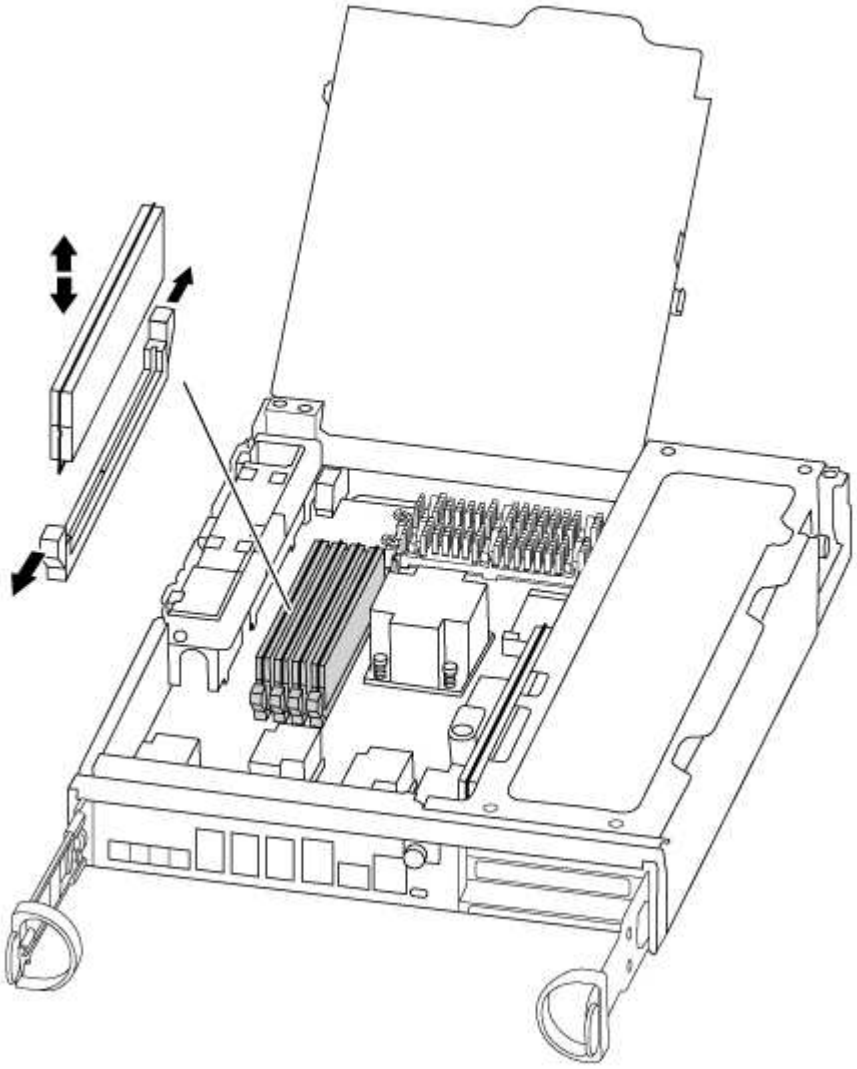

4. Locate the slot where you are installing the DIMM.
5. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Repeat these steps for the remaining DIMMs.
7. Move the NVMEM battery to the replacement controller module.
8. Align the tab or tabs on the battery holder with the notches in the controller module side, and then gently push down on the battery housing until the battery housing clicks into place.

Step 5: Move a PCIe card

To move PCIe cards, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

You must have the new controller module ready so that you can move the PCIe cards directly from the old controller module to the corresponding slots in the new one.

1. Loosen the thumbscrew on the controller module side panel.
2. Swing the side panel off the controller module.

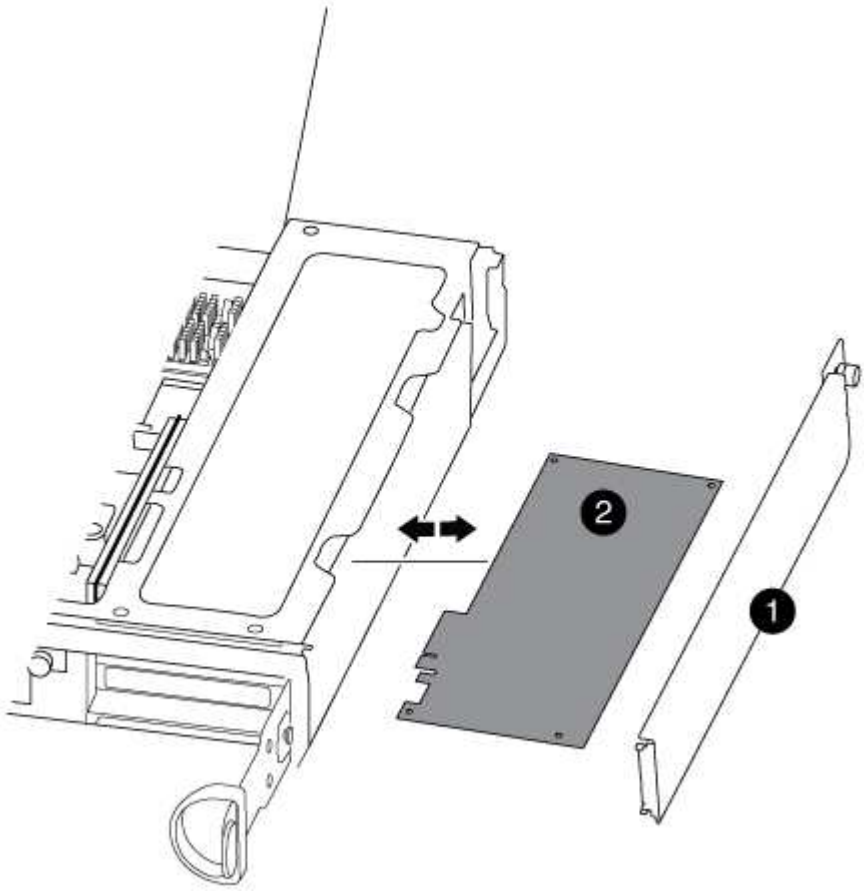

1	Side panel
2	PCIe card

3. Remove the PCIe card from the old controller module and set it aside.

   Make sure that you keep track of which slot the PCIe card was in.

4. Repeat the preceding step for the remaining PCIe cards in the old controller module.
5. Open the new controller module side panel, if necessary, slide off the PCIe card filler plate, as needed, and carefully install the PCIe card.

   Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in

the socket. The card must be fully and evenly seated in the slot.

6. Repeat the preceding step for the remaining PCIe cards that you set aside.
7. Close the side panel and tighten the thumbscrew.

Step 6: Move a caching module

You must move the caching modules from the impaired controller modules to the replacement controller module when replacing a controller module.

1. Locate the caching module at the rear of the controller module and remove it:
   a. Press the release tab.
   b. Remove the heatsink.

      The storage system comes with two slots available for the caching module and only one slot is occupied, by default.

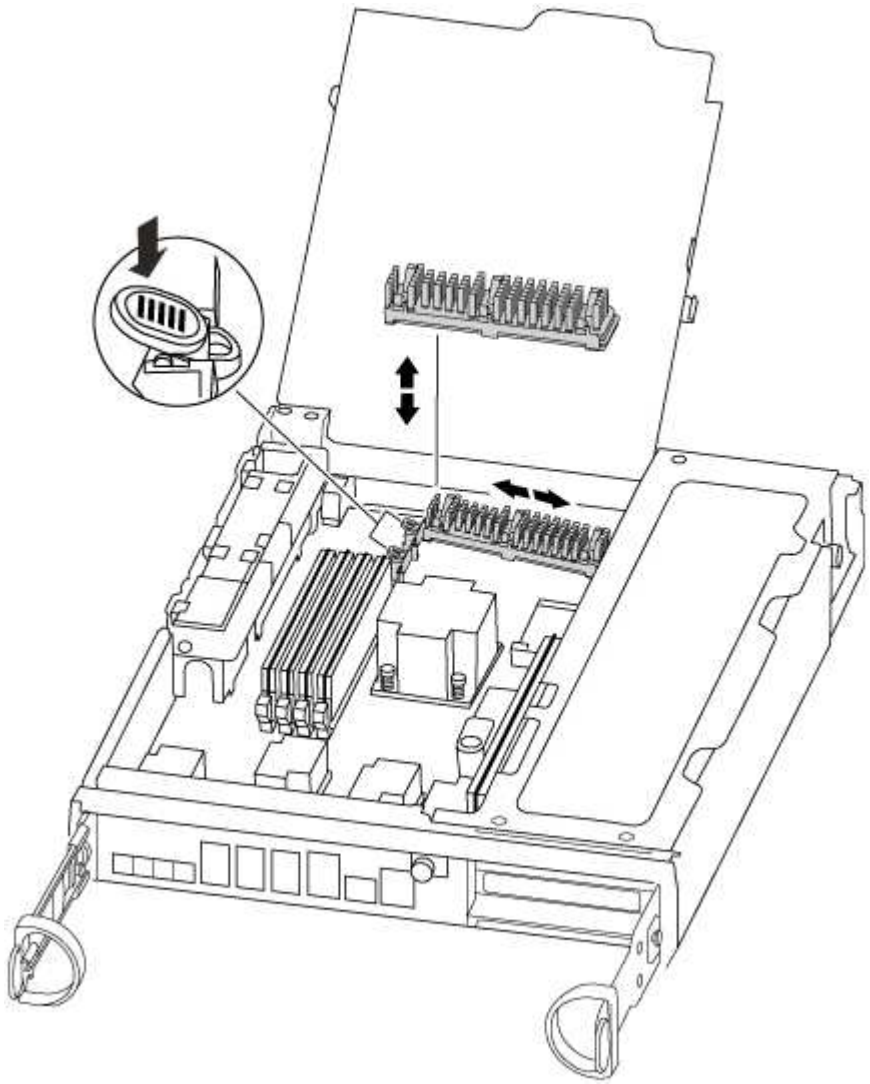

2. Move the caching module to the new controller module, and then align the edges of the caching module with the socket housing and gently push it into the socket.
3. Verify that the caching module is seated squarely and completely in the socket. If necessary, remove the caching module and reseat it into the socket.

4. Reseat and push the heatsink down to engage the locking button on the caching module housing.
5. Repeat the steps if you have a second caching module. Close the controller module cover.

Step 7: Install the controller

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

> The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, close the CPU air duct.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   > Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   > You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	`The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.` a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. When you see the message `Press Ctrl-C for Boot Menu`, press `Ctrl-C` to interrupt the boot process. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. Select the option to boot to Maintenance mode from the displayed menu.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process, and then press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. From the boot menu, select the option for Maintenance mode.

Important: During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down.
  You can safely respond `y` to these prompts.

Restore and verify the system configuration - FAS8200

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - FAS8200

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must use the correct procedure for your configuration.

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch. `boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node              Partner           Possible     State Description
------------      ------------      --------
-------------------------------------
node1             node2             false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2             node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys

   You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   a. Save any coredumps: `system node run -node local-node-name partner savecore`
   b. Wait for `savecore` command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the `savecore` command: `system node run -node local-node-name partner savecore -s`

   c. Return to the admin privilege level: `set -privilege admin`
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration content for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.
7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Option 2: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `` `metrocluster node show -fields node-systemid`,dr-partner-systemid` ``

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

```
dr-group-id cluster         node                 node-systemid dr-
partner-systemid
 ----------- -------------------- ------------------- -------------
------------------
 1           Cluster_A            Node_A_1             536872914
118073209
 1           Cluster_B            Node_B_1             118073209
536872914
 2 entries were displayed.
```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

```
Local System ID: 118065481
    ...
    ...
```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                POOL   SERIAL NUMBER  HOME
-------    -------------        -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:

a. Change to the advanced privilege level: `set -privilege advanced`

   You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

b. Verify that the coredumps are saved: `system node run -node *local-node-name* partner savecore`

   If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node *local-node-name* partner savecore -s command`.</info>.

c. Return to the admin privilege level: `set -privilege admin`

7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
8. Boot the *replacement* node: `boot_ontap`
9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`
10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`

c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - FAS8200

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`

b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS8200

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Open the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

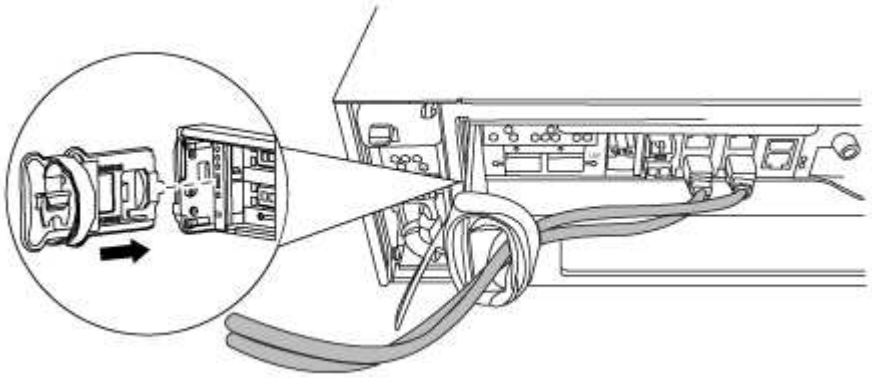

4. Loosen the thumbscrew on the cam handle on the controller module.

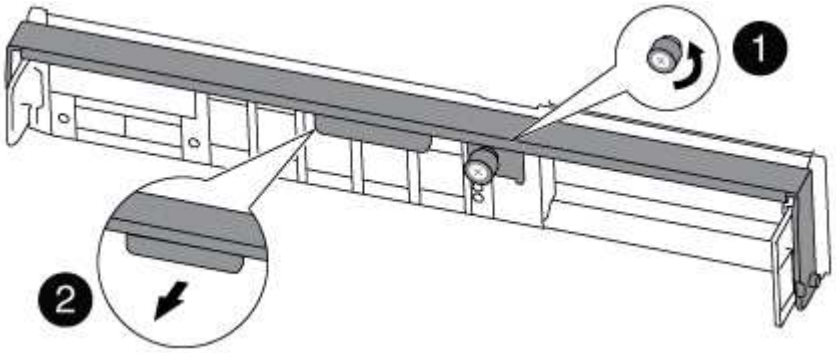

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

1. Check the NVMEM LED on the controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

   

2. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
3. Unplug the battery:

   The NVMEM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after Data ONTAP has successfully booted.

     a. Open the CPU air duct and locate the NVMEM battery.

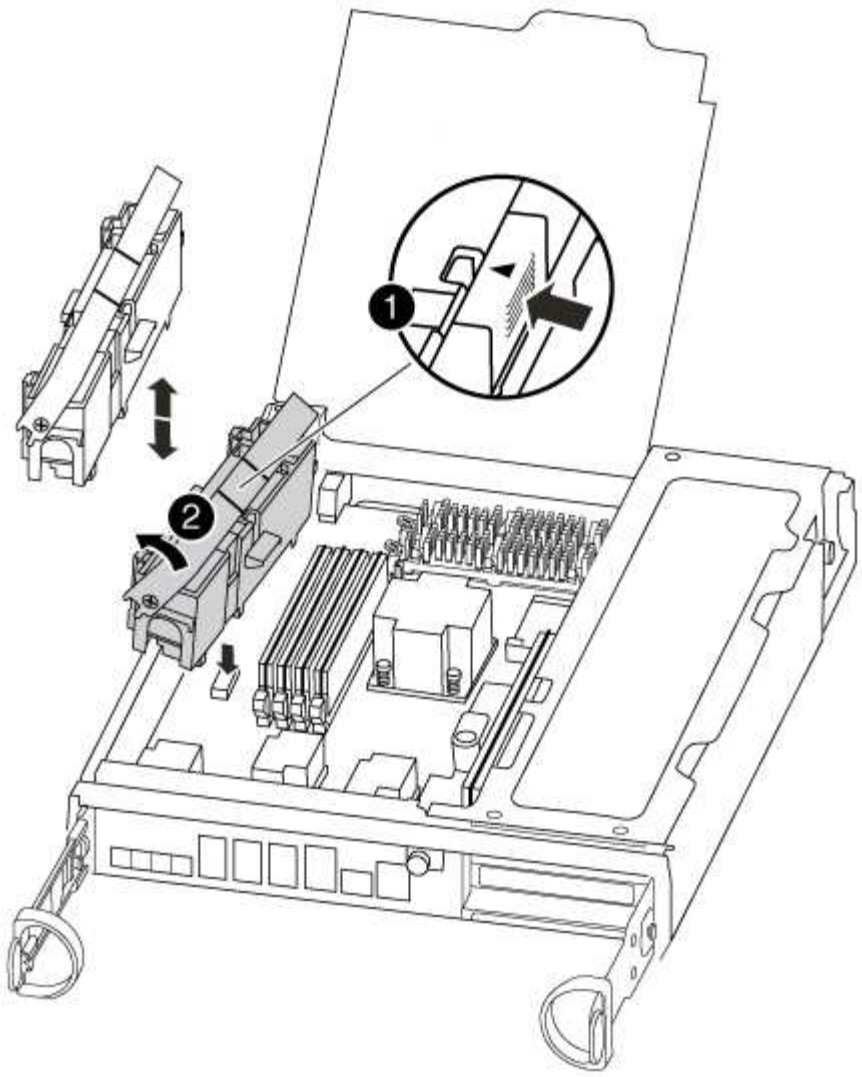

1	NVMEM battery lock tab
2	NVMEM battery

b. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.

c. Wait a few seconds, and then plug the battery back into the socket.

4. Check the NVMEM LED on the controller module.
5. Locate the DIMMs on your controller module.
6. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
7. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:

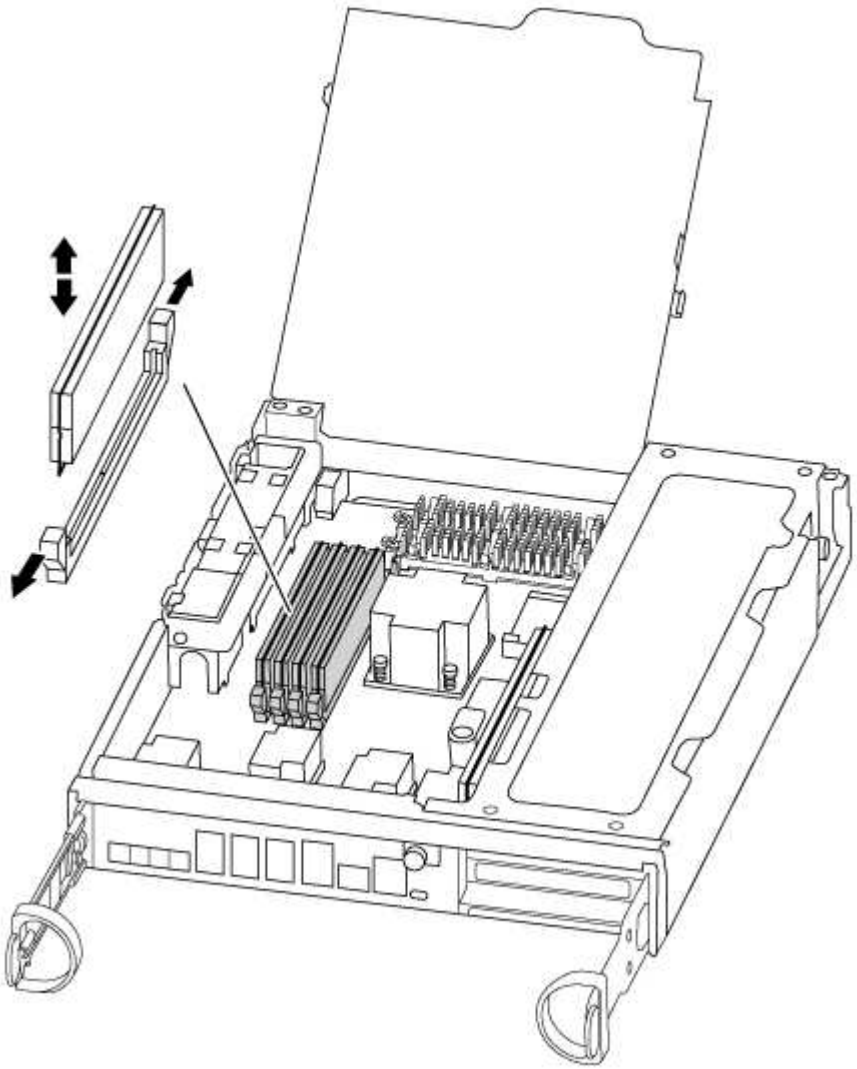

8. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

9. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

10. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
11. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

    Make sure that the plug locks down onto the controller module.

12. Close the controller module cover.

Step 4: Reinstall the controller

After you replace a component within the controller module, you must reinstall the controller module in the system chassis.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

3. Complete the reinstallation of the controller module:

   The controller module begins to boot as soon as it is fully seated in the chassis.

   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Tighten the thumbscrew on the cam handle on back of the controller module.

   c. If you have not already done so, reinstall the cable management device.

   d. Bind the cables to the cable management device with the hook and loop strap.

Step 5 (Two-node MetroCluster only): Switch back aggregates

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a fan - FAS8200

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press down the release latch on the fan module cam handle, and then pull the cam handle downward.

   The fan module moves a little bit away from the chassis.

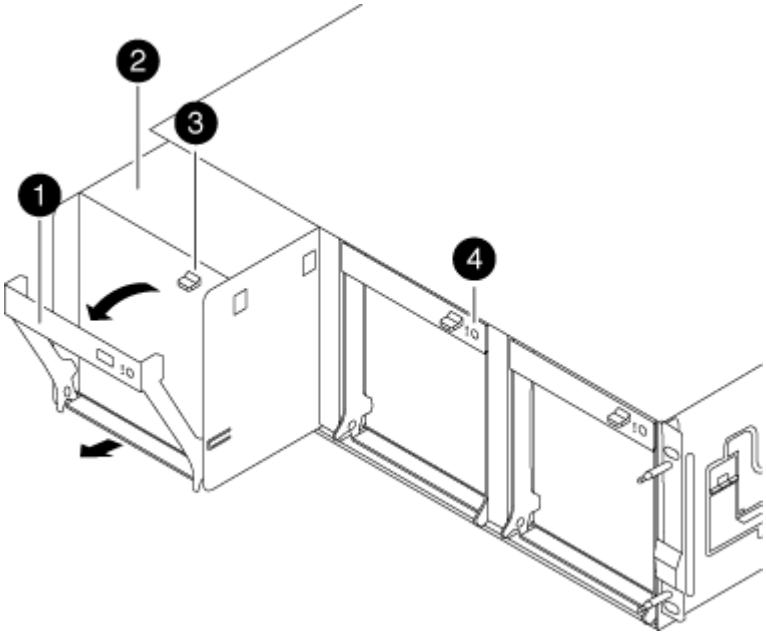

1	Cam handle
2	Fan module
3	Cam handle release latch

4	Fan module Attention LED

5. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

6. Set the fan module aside.
7. Insert the replacement fan module into the chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The fan LED should be green after the fan is seated and has spun up to operational speed.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVMEM battery - FAS8200

To replace an NVMEM battery in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Open the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

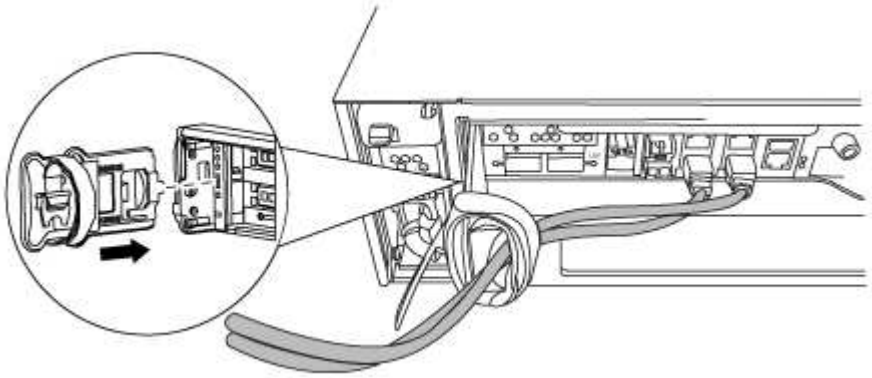

4. Loosen the thumbscrew on the cam handle on the controller module.

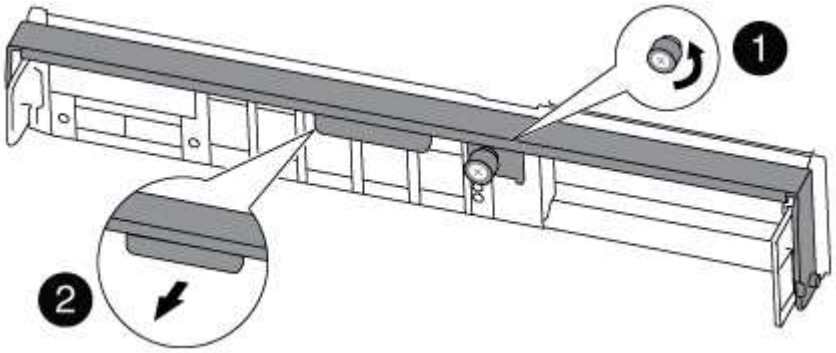

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery in your system, you must remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

> The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Open the CPU air duct and locate the NVMEM battery.

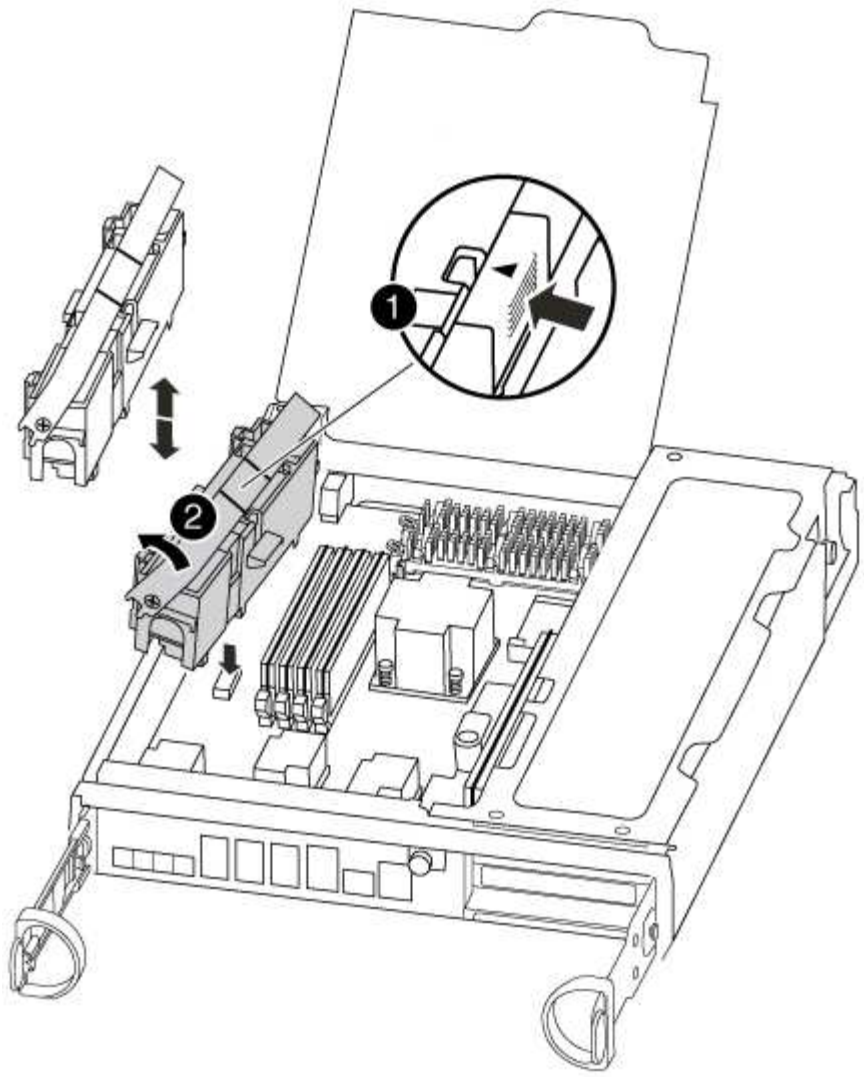

1	Battery lock tab
2	NVMEM battery pack

3. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
4. Remove the replacement battery from its package.
5. Align the tab or tabs on the battery holder with the notches in the controller module side, and then gently push down on the battery housing until the battery housing clicks into place.
6. Close the CPU air duct.

   Make sure that the plug locks down to the socket.

Step 4: Reinstall the controller

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

3. Complete the reinstallation of the controller module:

   The controller module begins to boot as soon as it is fully seated in the chassis.

   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Tighten the thumbscrew on the cam handle on back of the controller module.

   c. If you have not already done so, reinstall the cable management device.

   d. Bind the cables to the cable management device with the hook and loop strap.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`

4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe card - FAS8200

To replace a PCIe card, you must perform a specific sequence of tasks.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

   If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Open the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

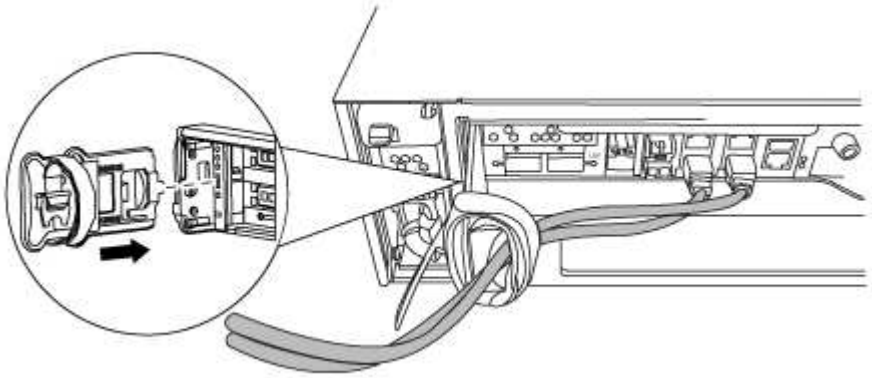

4. Loosen the thumbscrew on the cam handle on the controller module.

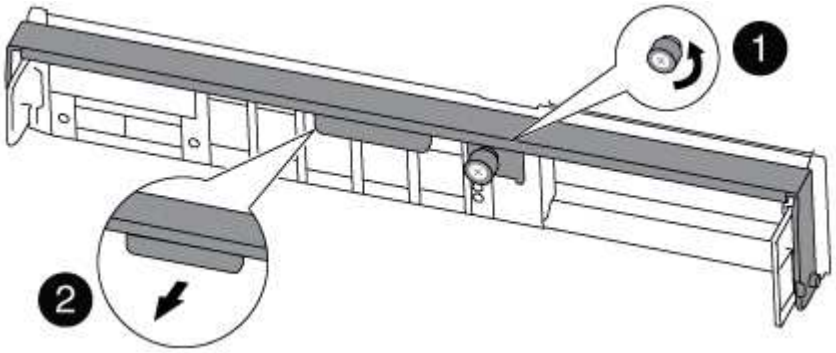

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace a PCIe card

To replace a PCIe card, locate it within the controller and follow the specific sequence of steps.

1. Loosen the thumbscrew on the controller module side panel.
2. Swing the side panel off the controller module.

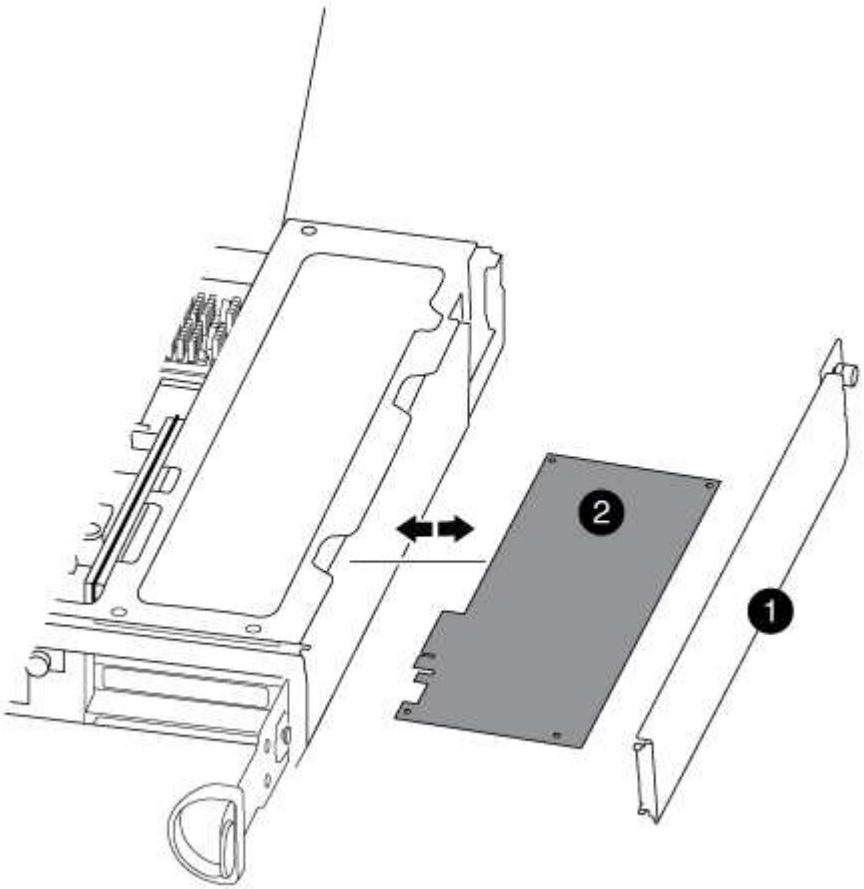

1	Side panel
2	PCIe card

3. Remove the PCIe card from the controller module and set it aside.
4. Install the replacement PCIe card.

   Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The PCIe card must be fully and evenly seated in the slot.

   If you are installing a card in the bottom slot and cannot see the card socket well, remove the top card so that you can see the card socket, install the card, and then reinstall the card you removed from the top slot.

5. Close the side panel and tighten the thumbscrew.

Step 4: Reinstall the controller

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

3. Complete the reinstallation of the controller module:

   The controller module begins to boot as soon as it is fully seated in the chassis.

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. If you have not already done so, reconnect the cables to the controller module. d. Bind the cables to the cable management device with the hook and loop strap.
A two-node MetroCluster configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Tighten the thumbscrew on the cam handle on back of the controller module. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. If you have not already done so, reconnect the cables to the controller module. d. Bind the cables to the cable management device with the hook and loop strap. e. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.

4. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs or onboard ports, convert these ports to 10 GbE connections by using the nicadmin convert command from Maintenance mode.

   Be sure to exit Maintenance mode after completing the conversion.

5. Return the controller to normal operation:

If your system is in…	Issue this command from the partner's console…
An HA pair	`storage failover giveback -ofnode` *`impaired_node_name`*
A two-node MetroCluster configuration	Proceed to the next step. The MetroCluster switchback procedure is done in the next task in the replacement process.

6. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5 (two-node MetroCluster only): Switch back aggregate

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - FAS8200

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- The number of power supplies in the system depends on the model.
- Power supplies are auto-ranging.

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.

3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Press down the release latch on the power supply cam handle, and then lower the cam handle to the fully open position to release the power supply from the mid plane.

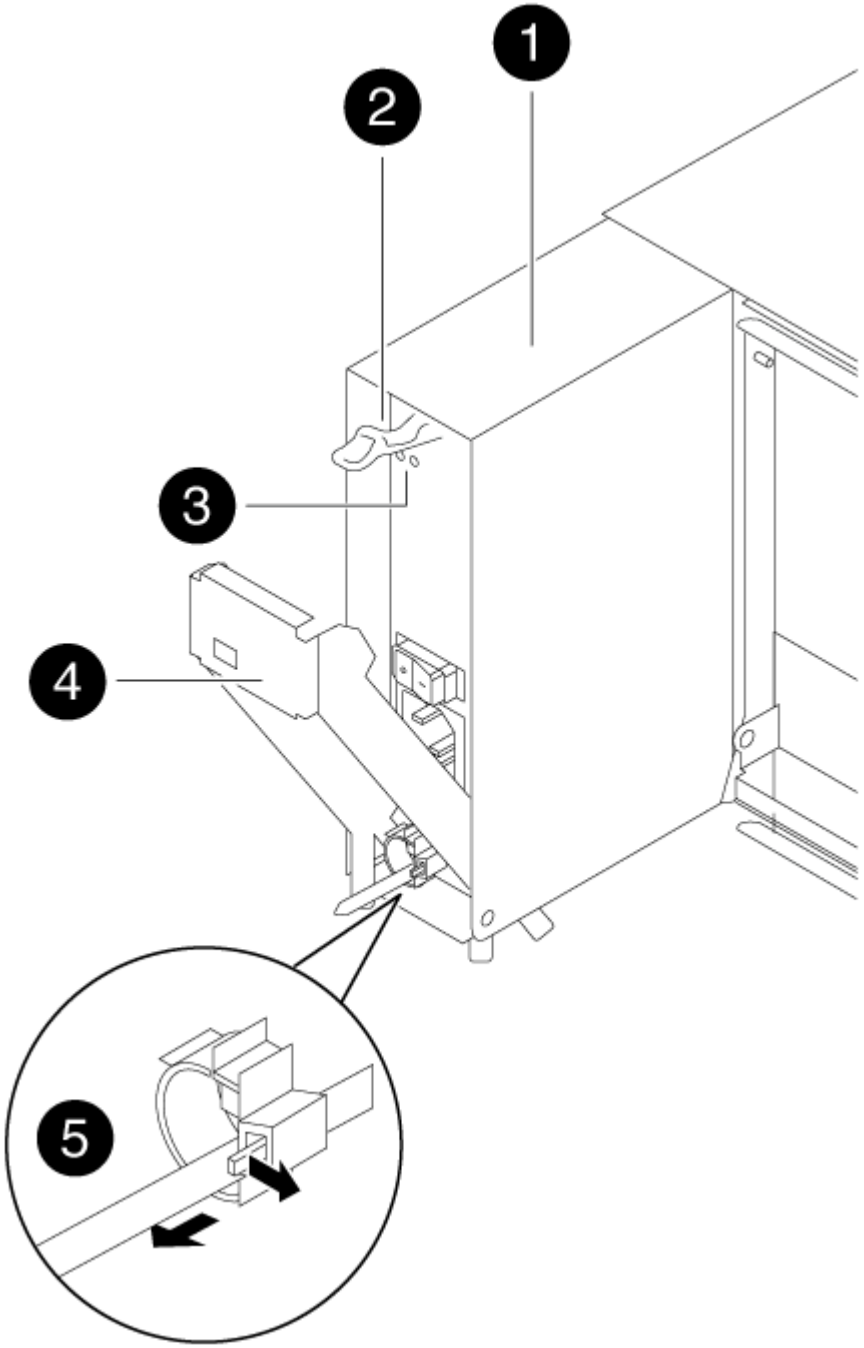

1	Power supply

2	Cam handle release latch
3	Power and Fault LEDs
4	Cam handle
5	Power cable locking mechanism

5. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.
7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Push firmly on the power supply cam handle to seat it all the way into the chassis, and then push the cam handle to the closed position, making sure that the cam handle release latch clicks into its locked position.
9. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

Once power is restored to the power supply, the status LED should be green.

1. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

   The power supply LEDs are lit when the power supply comes online.

2. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - FAS8200

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Open the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove and set aside the cable management devices from the left and right sides of the controller module.

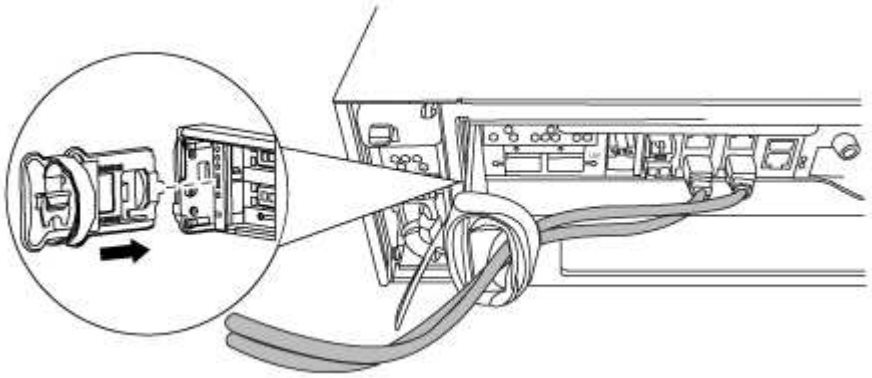

4. Loosen the thumbscrew on the cam handle on the controller module.

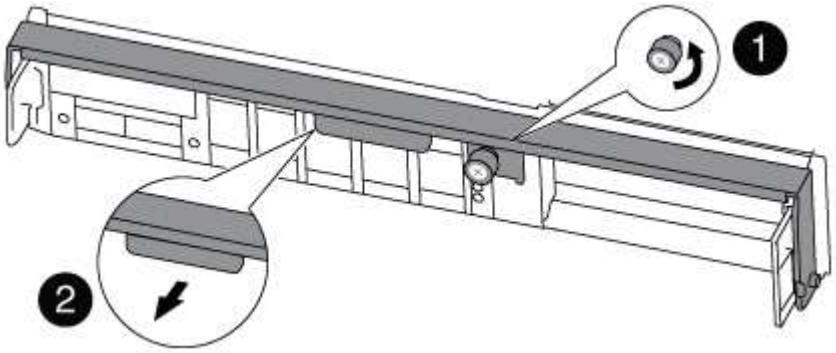

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

Step 3: Replace the RTC Battery

To replace the RTC battery, locate them inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

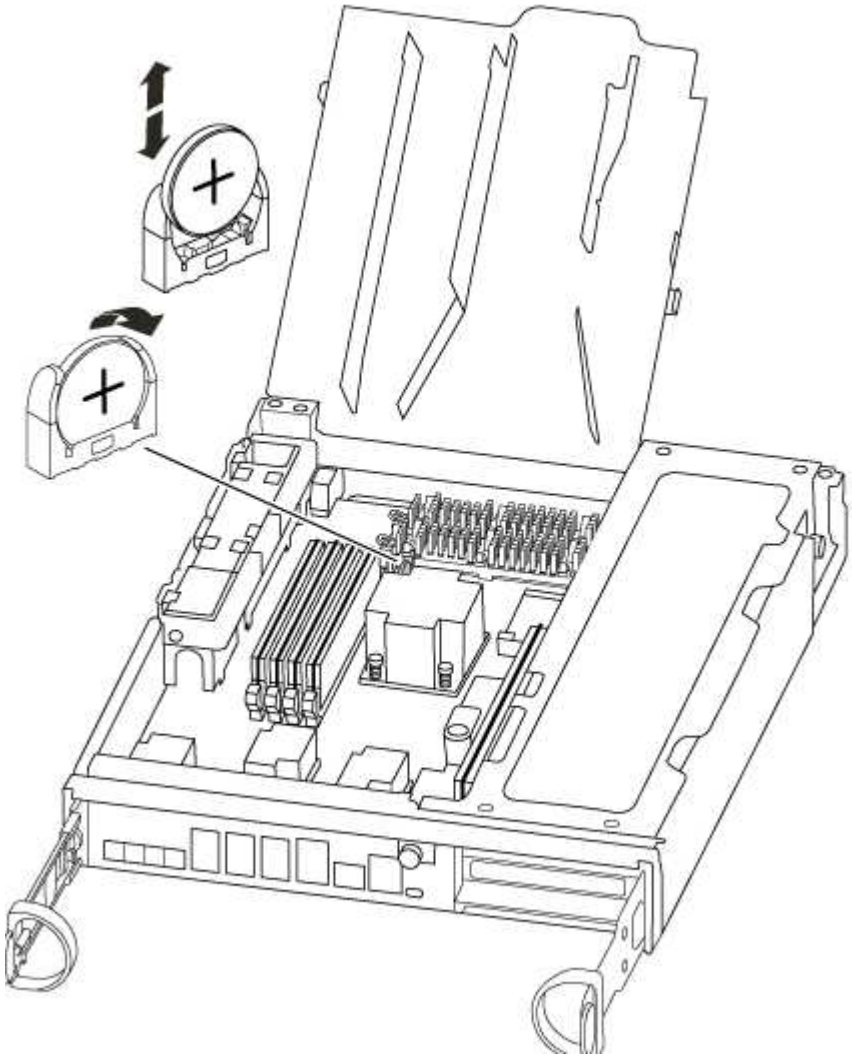

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module and setting time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Tighten the thumbscrew on the cam handle on back of the controller module.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

FAS9000 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

For MetroCluster configurations, see either:

- Install MetroCluster IP configuration
- Install MetroCluster Fabric-Attached configuration

Quick steps - FAS9000

This guide gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Access the *Installation and Setup Instructions* PDF poster:

AFF A700 Installation and Setup Instructions

FAS9000 Installation and Setup Instructions

Video steps - FAS9000

The following video shows how to install and cable your new system.

Animation - Install and setup of an AFF A700 or FAS9000

Detailed guide - FAS9000

This guide gives detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your system, you need to create an account on the NetApp Support Site, register your system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

You need to provide the following at your site:

- Rack space for the storage system
- Phillips #2 screwdriver
- Additional networking cables to connect your system to your network switch and laptop or console with a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.

3. Inventory and make a note of the number and types of cables you received.

   The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
10 GbE network cable	X6566B-2-R6, (112-00299), 2m X6566B-3-R6, 112-00300, 3m X6566B-5-R6 , 112-00301, 5m		Network cable

Type of cable…	Part number and length	Connector type	For…
40 GbE network cable 40 GbE cluster interconnect	X66100-1,112-00542, 1m X66100-3,112-00543, 3m		40 GbE network Cluster interconnect
100 GbE network cable 100 GbE storage cable	X66211A-05 (112-00595), 0.5m X66211A-1 (112-00573), 1m X66211A-2 (112-00574), 2m X66211A-5 (112-00574), 5m		Network cable Storage cable This cable applies to AFF A700 only.
Optical network cables (order dependent)	X6553-R6 (112-00188), 2m X6536-R6 (112-00090), 5m		FC host network
Cat 6, RJ-45 (order dependent)	Part numbers X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network and Ethernet data
Storage	X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		Storage
Micro-USB console cable	Not applicable		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	Not applicable		Powering up the system

4. Review the *NetApp ONTAP Configuration Guide* and collect the required information listed in that guide.

   ONTAP Configuration Guide

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

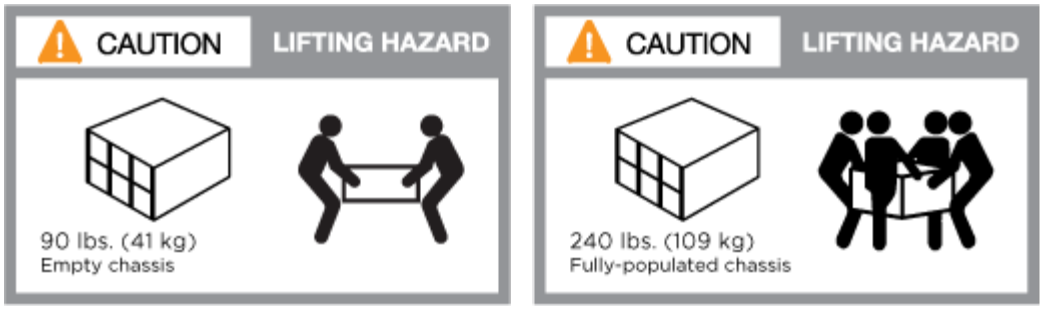

The label on the left indicates an empty chassis, while the label on the right indicates a fully-populated system.

3. Attach cable management devices (as shown).

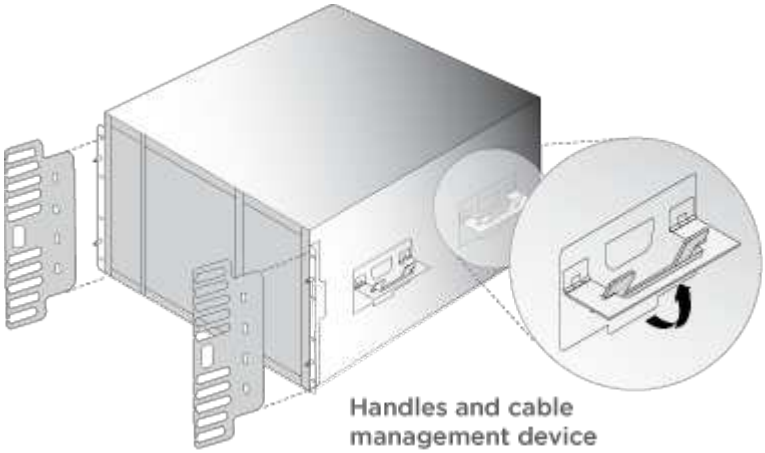

4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Two-node switchless cluster

Management network, data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all networking module ports.

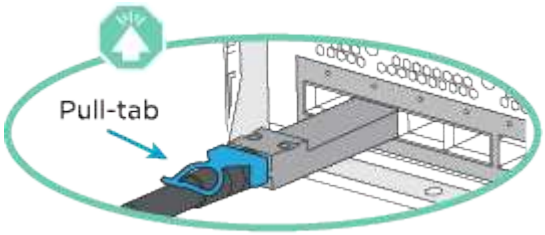

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Cable a two-node switchless cluster

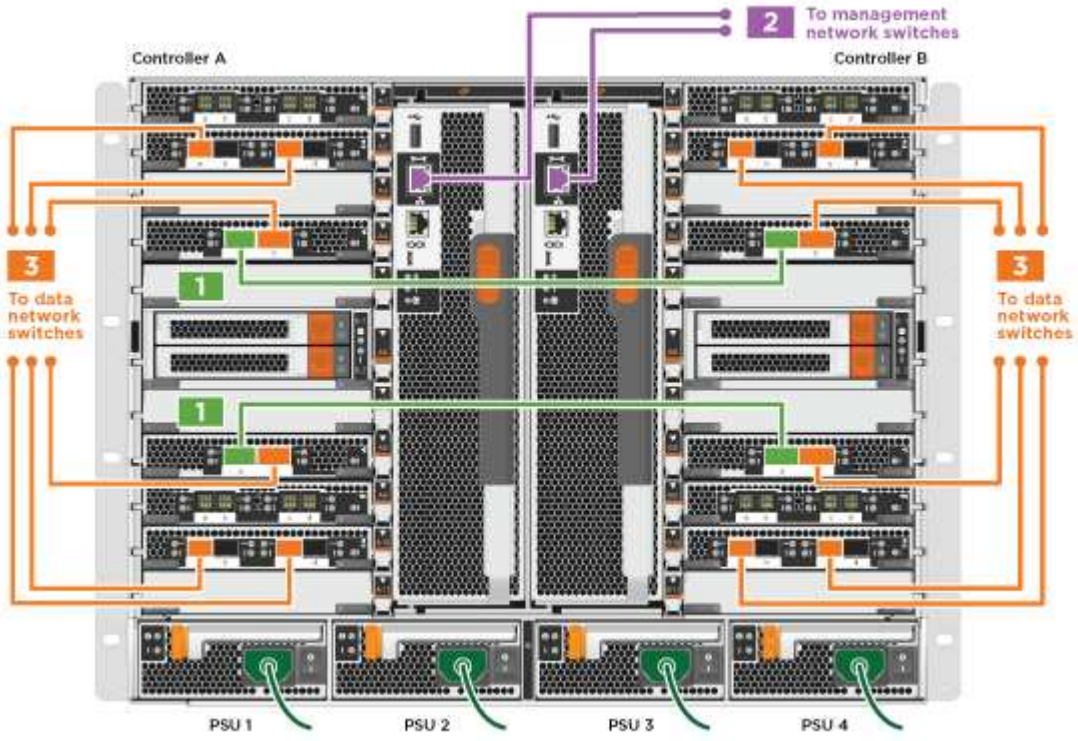

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Option 2: Switched cluster

Management network, data network, and management ports on the controllers are connected to switches. The cluster interconnect and HA ports are cabled on to the cluster/HA switch.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all networking module ports.

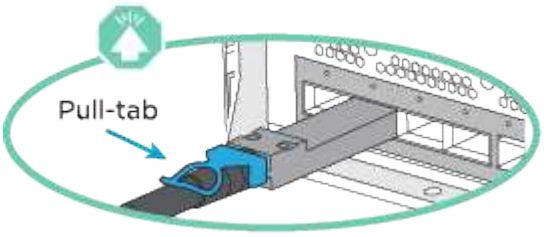

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Switched cluster cabling

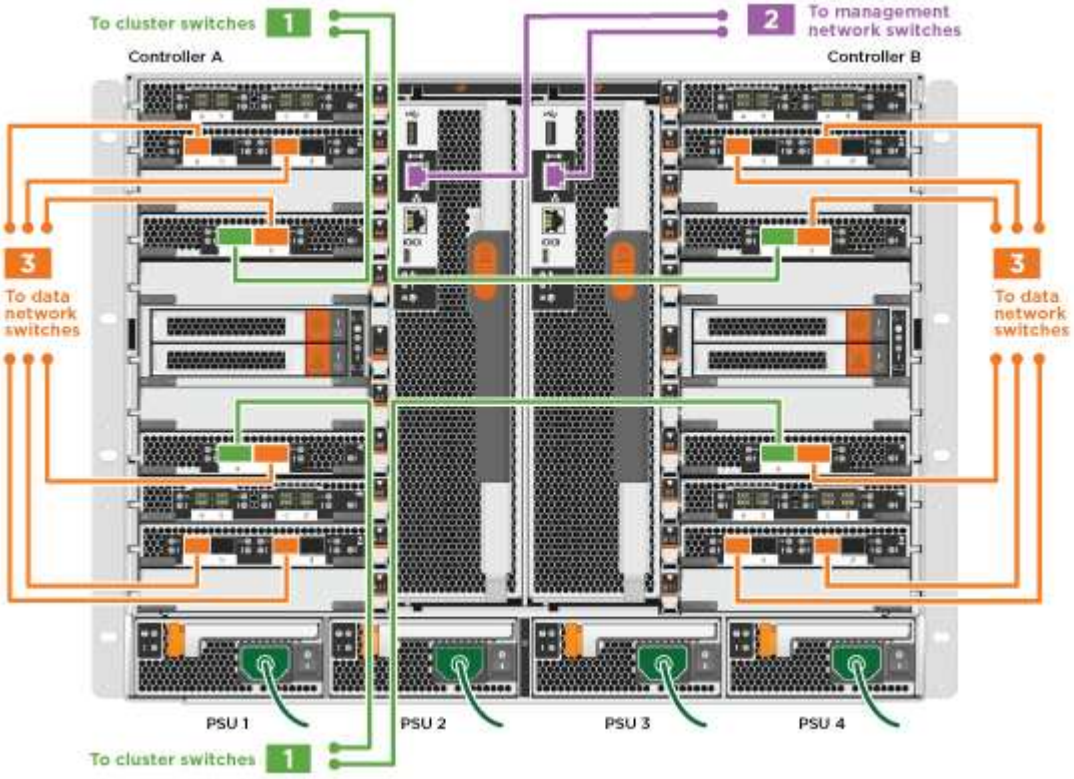

2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Step 4: Cable controllers to drive shelves

You can cable your new system to DS212C, DS224C, or NS224 shelves, depending on if it is an AFF or FAS system.

Option 1: Cable the controllers to DS212C or DS224C drive shelves

You must cable the shelf-to-shelf connections, and then cable both controllers to the DS212C or DS224C drive shelves.

The cables are inserted into the drive shelf with the pull-tabs facing down, while the other end of the cable is inserted into the controller storage modules with the pull-tabs up.

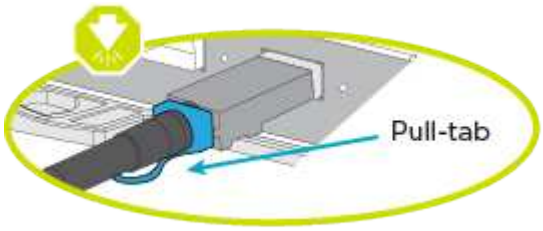

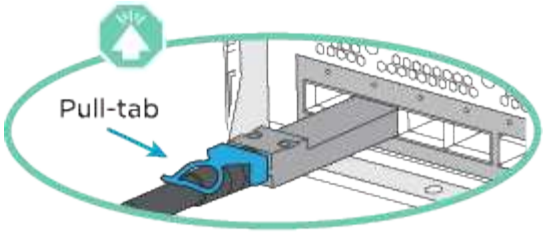

Steps

1. Use the following animations or illustrations to cable your drive shelves to your controllers.

   The examples use DS224C shelves. Cabling is similar with other supported SAS drive shelves.

   - Cabling SAS shelves in FAS9000, AFF A700, and ASA AFF A700, ONTAP 9.7 and earlier:

     Animation - Cable SAS storage - ONTAP 9.7 and earlier

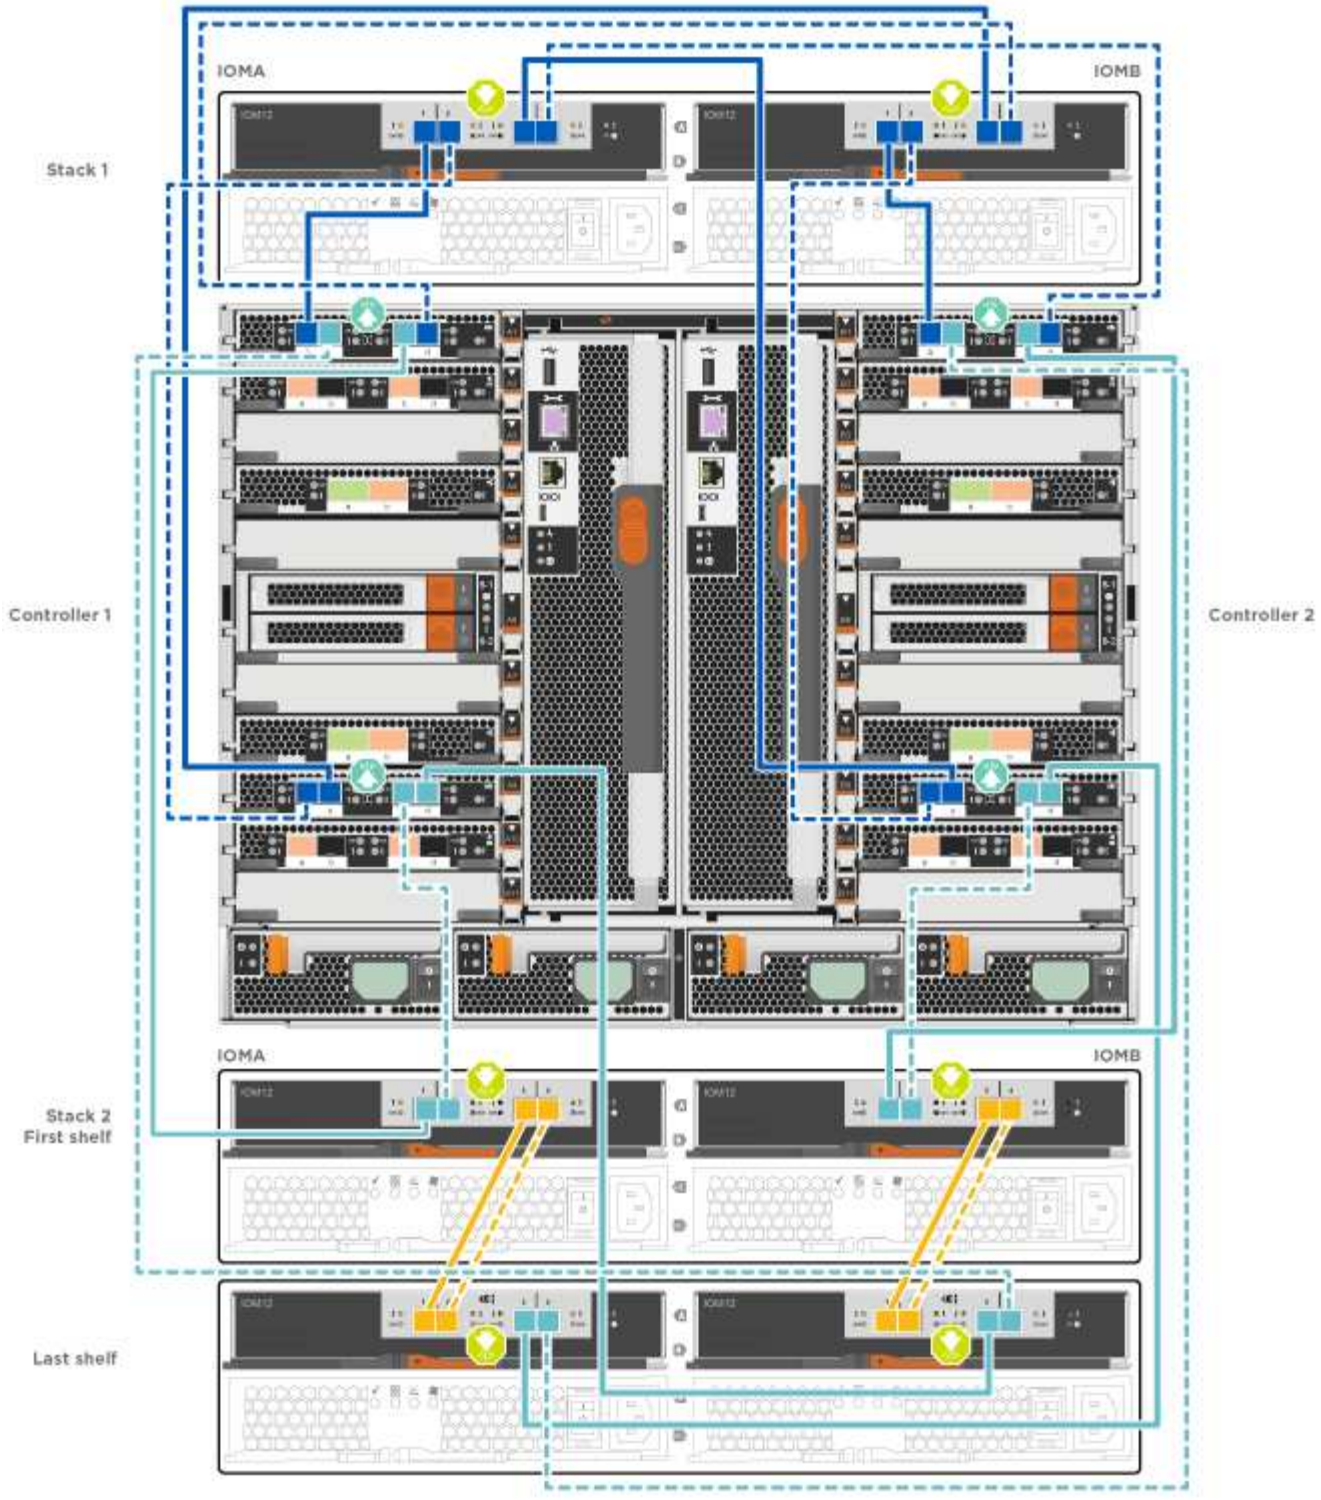

- Cabling SAS shelves in FAS9000, AFF A700, and ASA AFF A700, ONTAP 9.8 and later:

  Animation - Cable SAS storage - ONTAP 9.8 and later

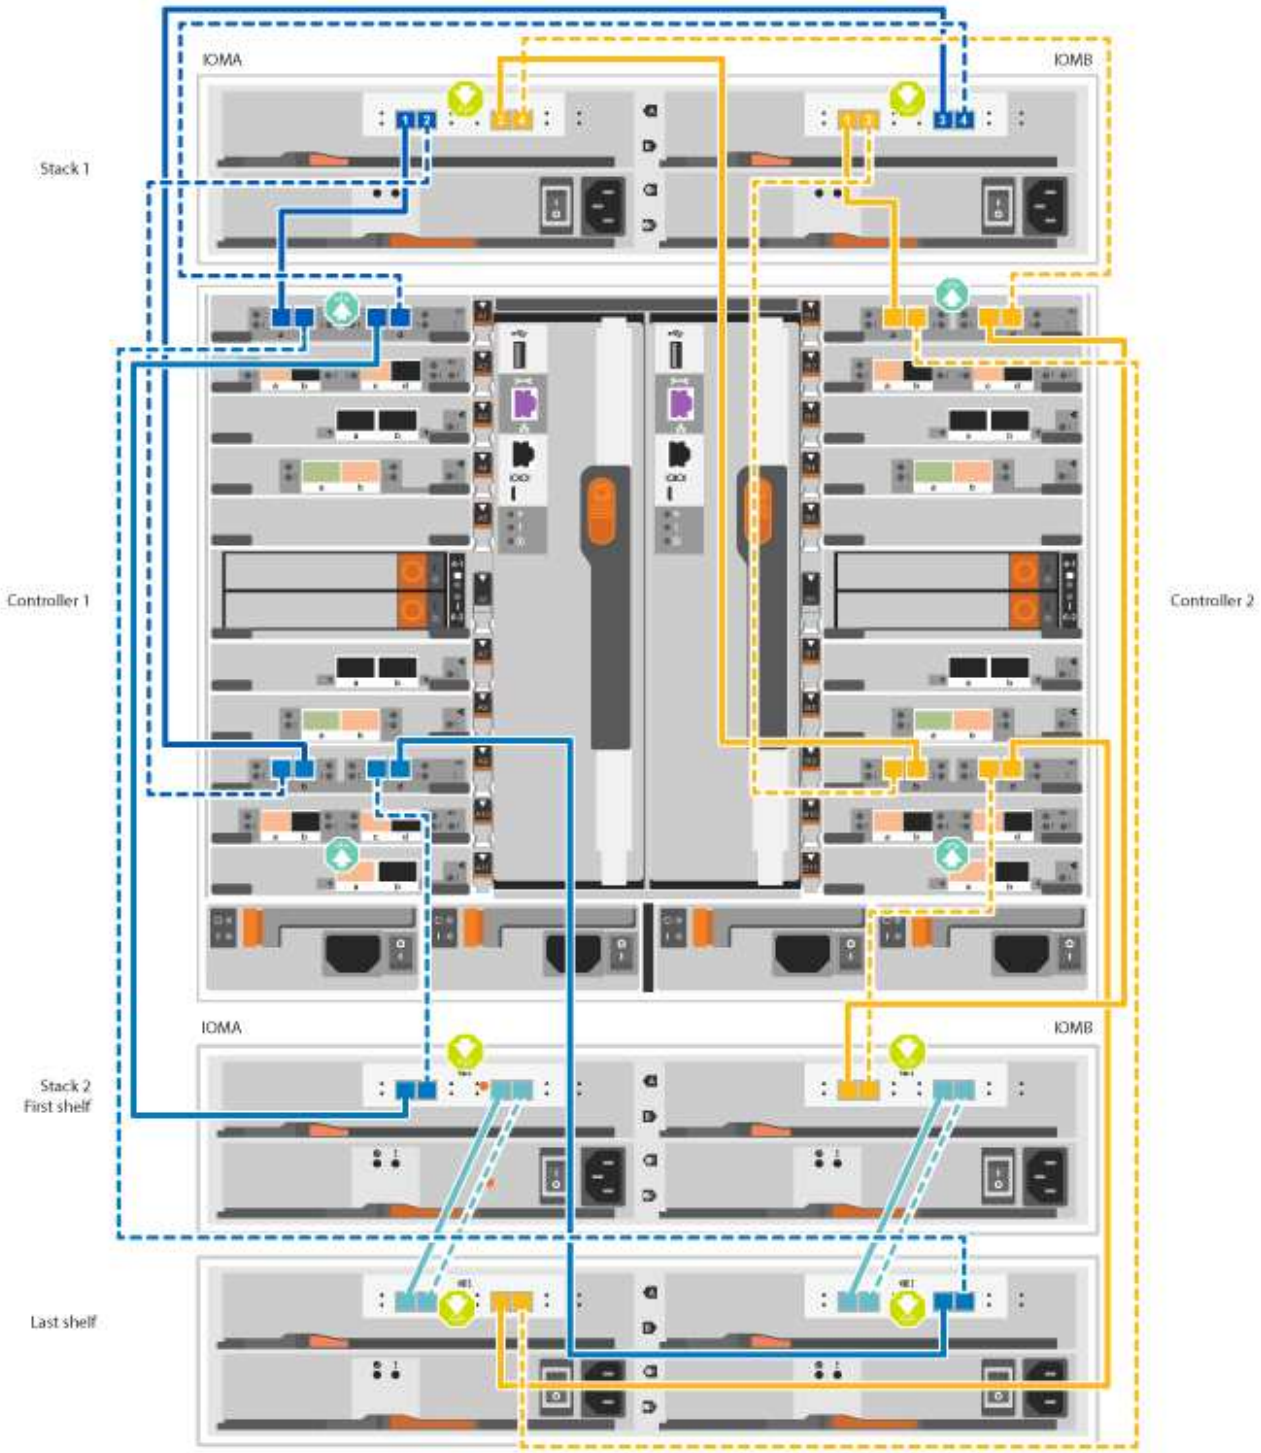

If you have more than one drive shelf stack, see the *Installation and Cabling Guide* for your drive shelf type.

Install and cable shelves for a new system installation - shelves with IOM12 modules

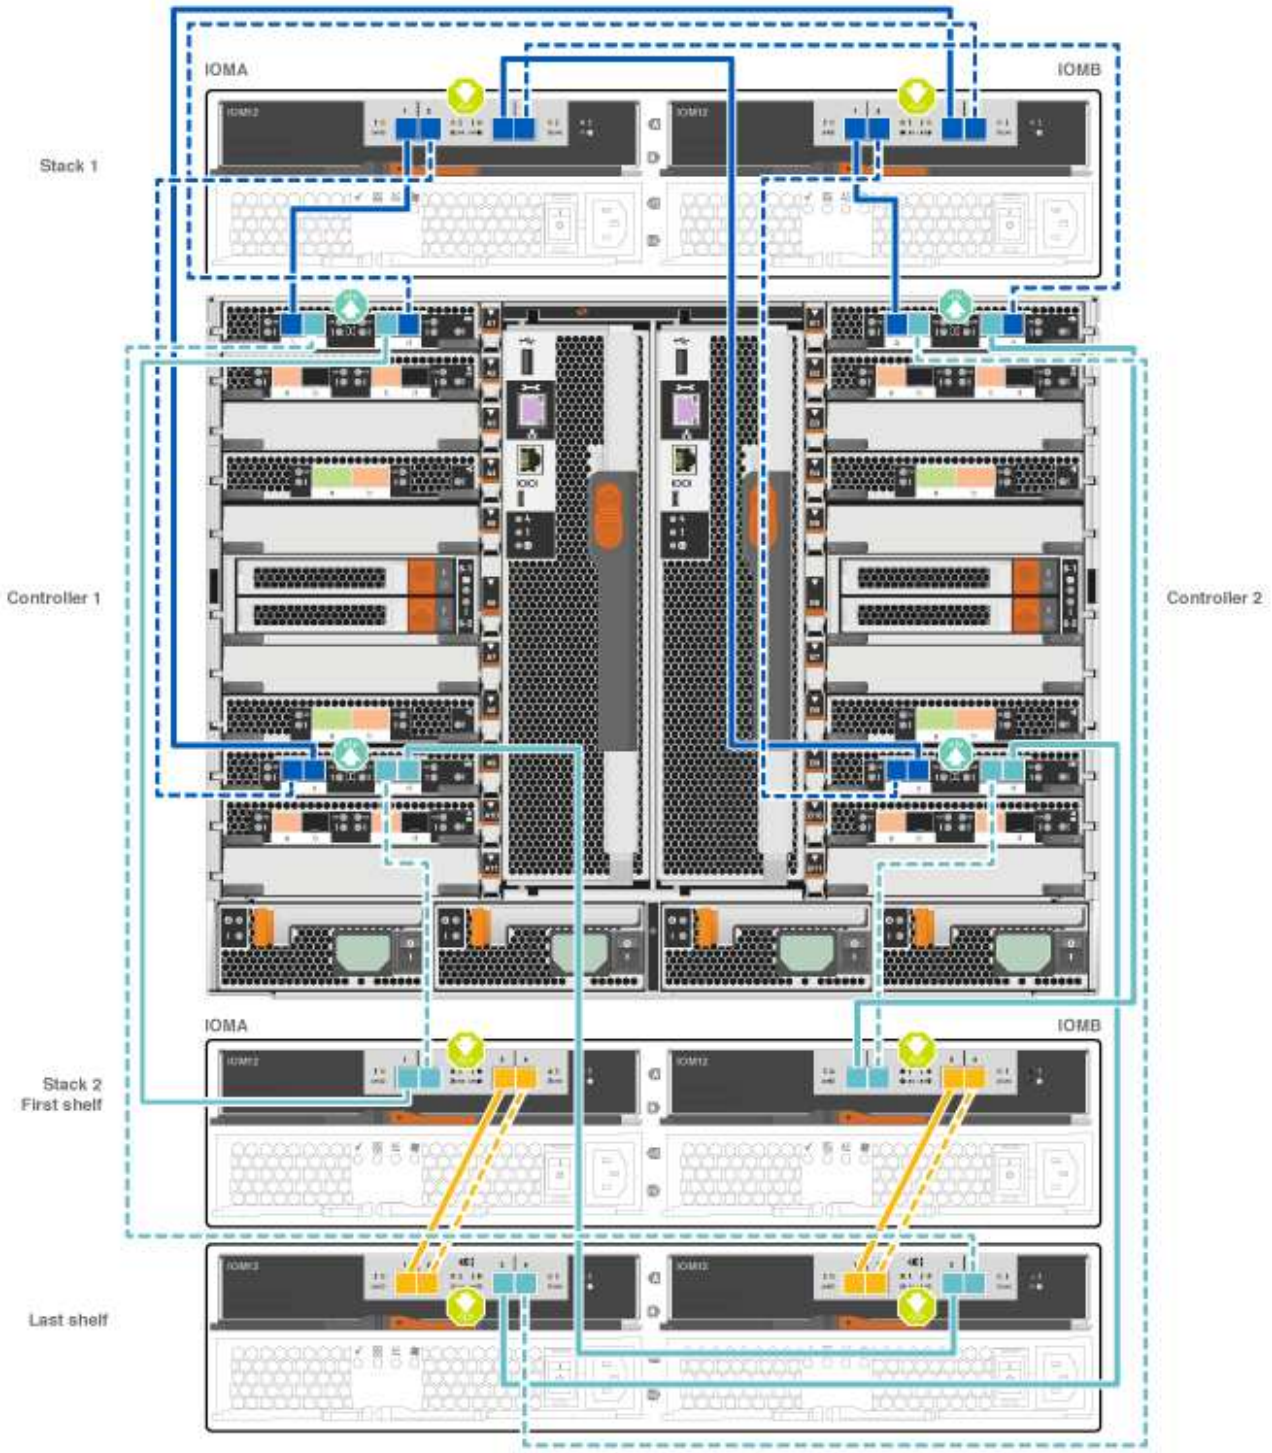

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Option 2: Cable the controllers to a single NS224 drive shelf in ASA AFF A700 systems running ONTAP 9.8 and later only

You must cable each controller to the NSM modules on the NS224 drive shelf on an AFF A700 or ASA AFF A700 running system ONTAP 9.8 or later.

- This task applies to ASA AFF A700 running ONTAP 9.8 or later only.
- The systems must have at least one X91148A module installed in slots 3 and/or 7 for each controller. The

animation or illustrations show this module installed in both slots 3 and 7.

- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the storage modules are up, while the pull tabs on the shelves are down.

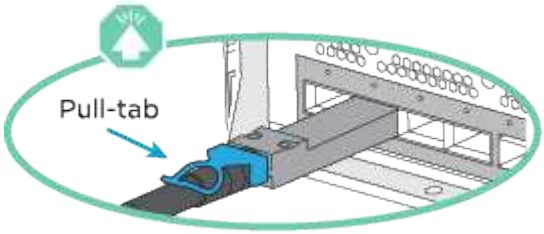

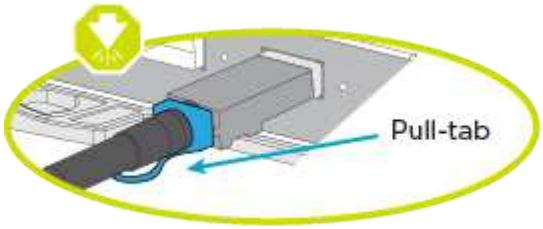

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following animation or illustrations to cable your controllers with two X91148A storage modules to a single NS224 drive shelf, or use the diagram to cable your controllers with one X91148A storage module to a single NS224 drive shelf.

   Animation - Cable a single NS224 shelf - ONTAP 9.8 and later

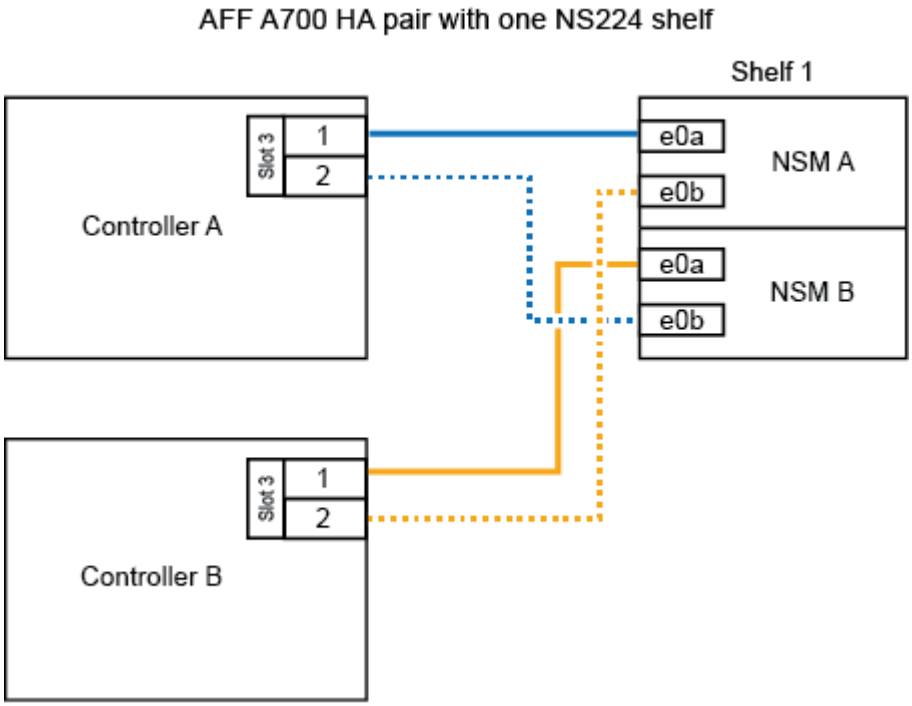

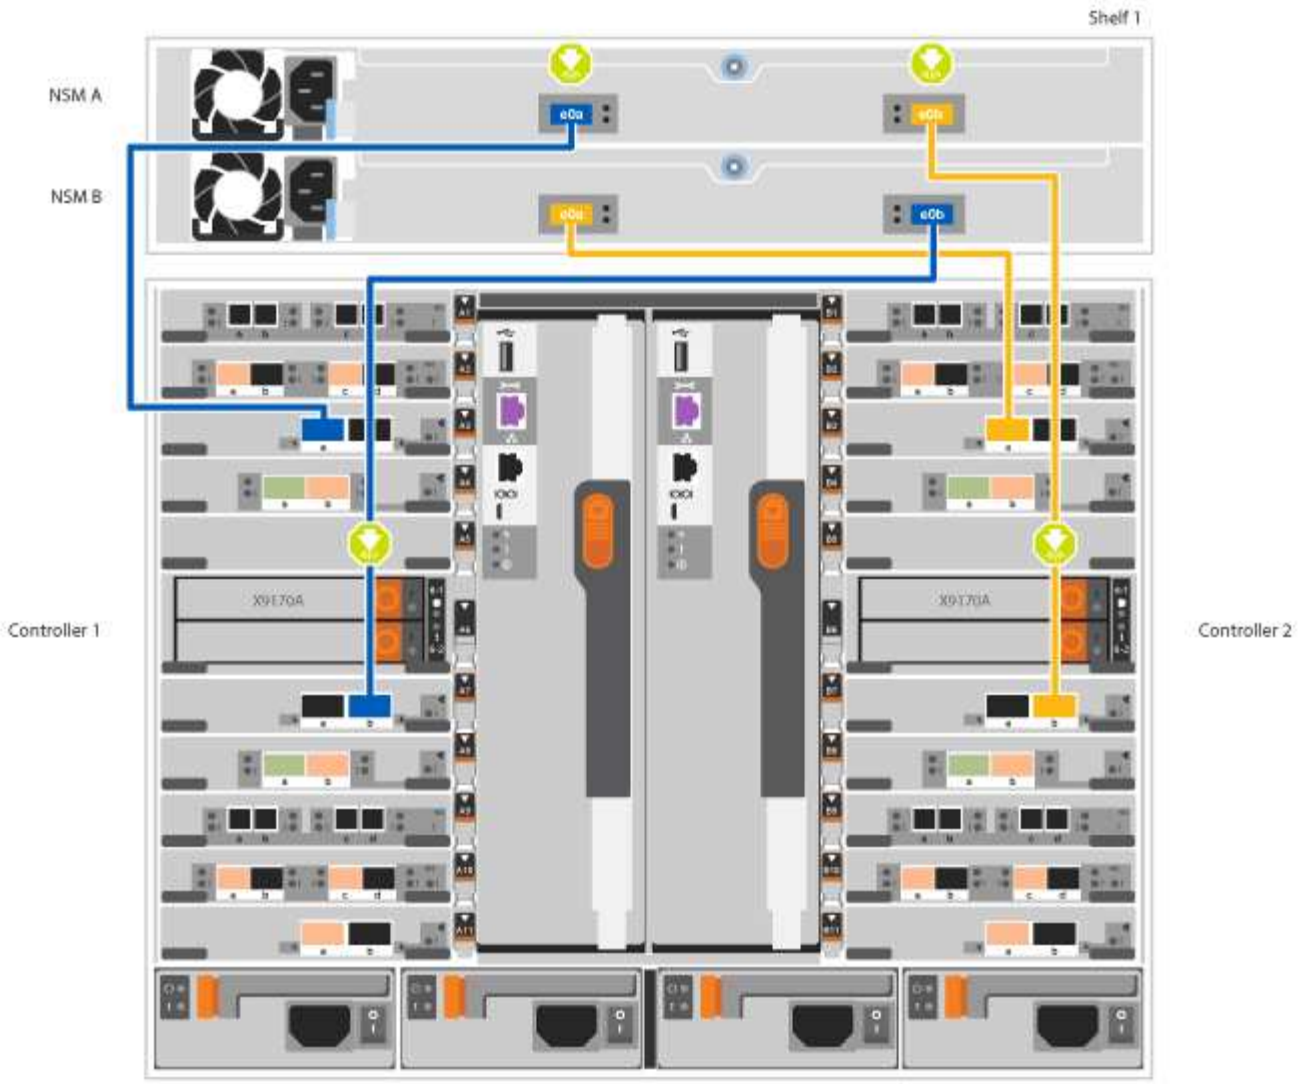

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Option 3: Cable the controllers to two NS224 drive shelves in ASA AFF A700 systems running ONTAP 9.8 and later only

You must cable each controller to the NSM modules on the NS224 drive shelves on an AFF A700 or ASA AFF A700 running system ONTAP 9.8 or later.

- This task applies to ASA AFF A700 running ONTAP 9.8 or later only.
- The systems must have two X91148A modules, per controller, installed in slots 3 and 7.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the storage modules are up, while the pull tabs on the shelves are down.

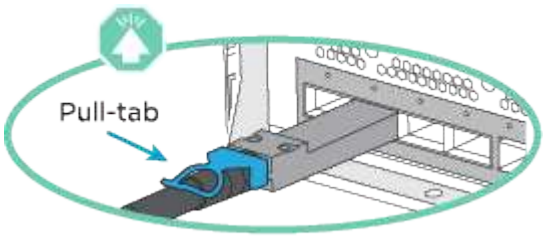

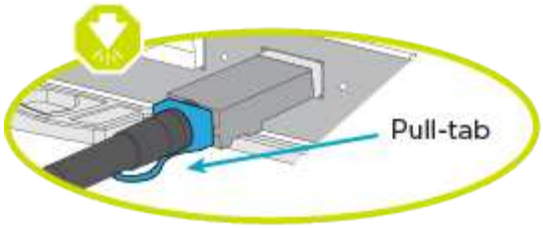

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following animation or illustrations to cable your controllers to two NS224 drive shelves.

Animation - Cable two NS224 shelves - ONTAP 9.8 and later

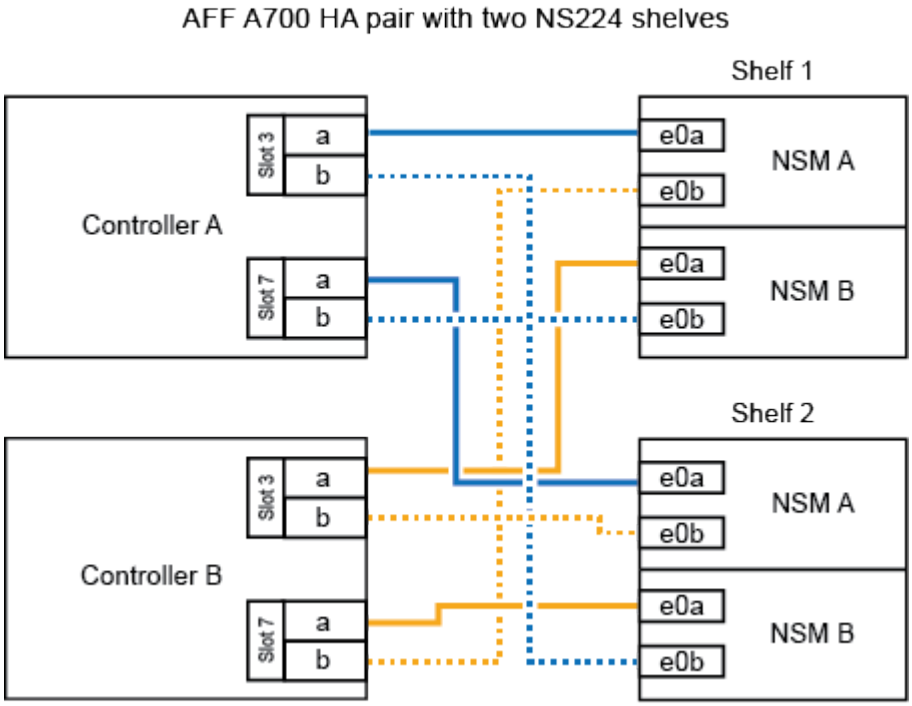

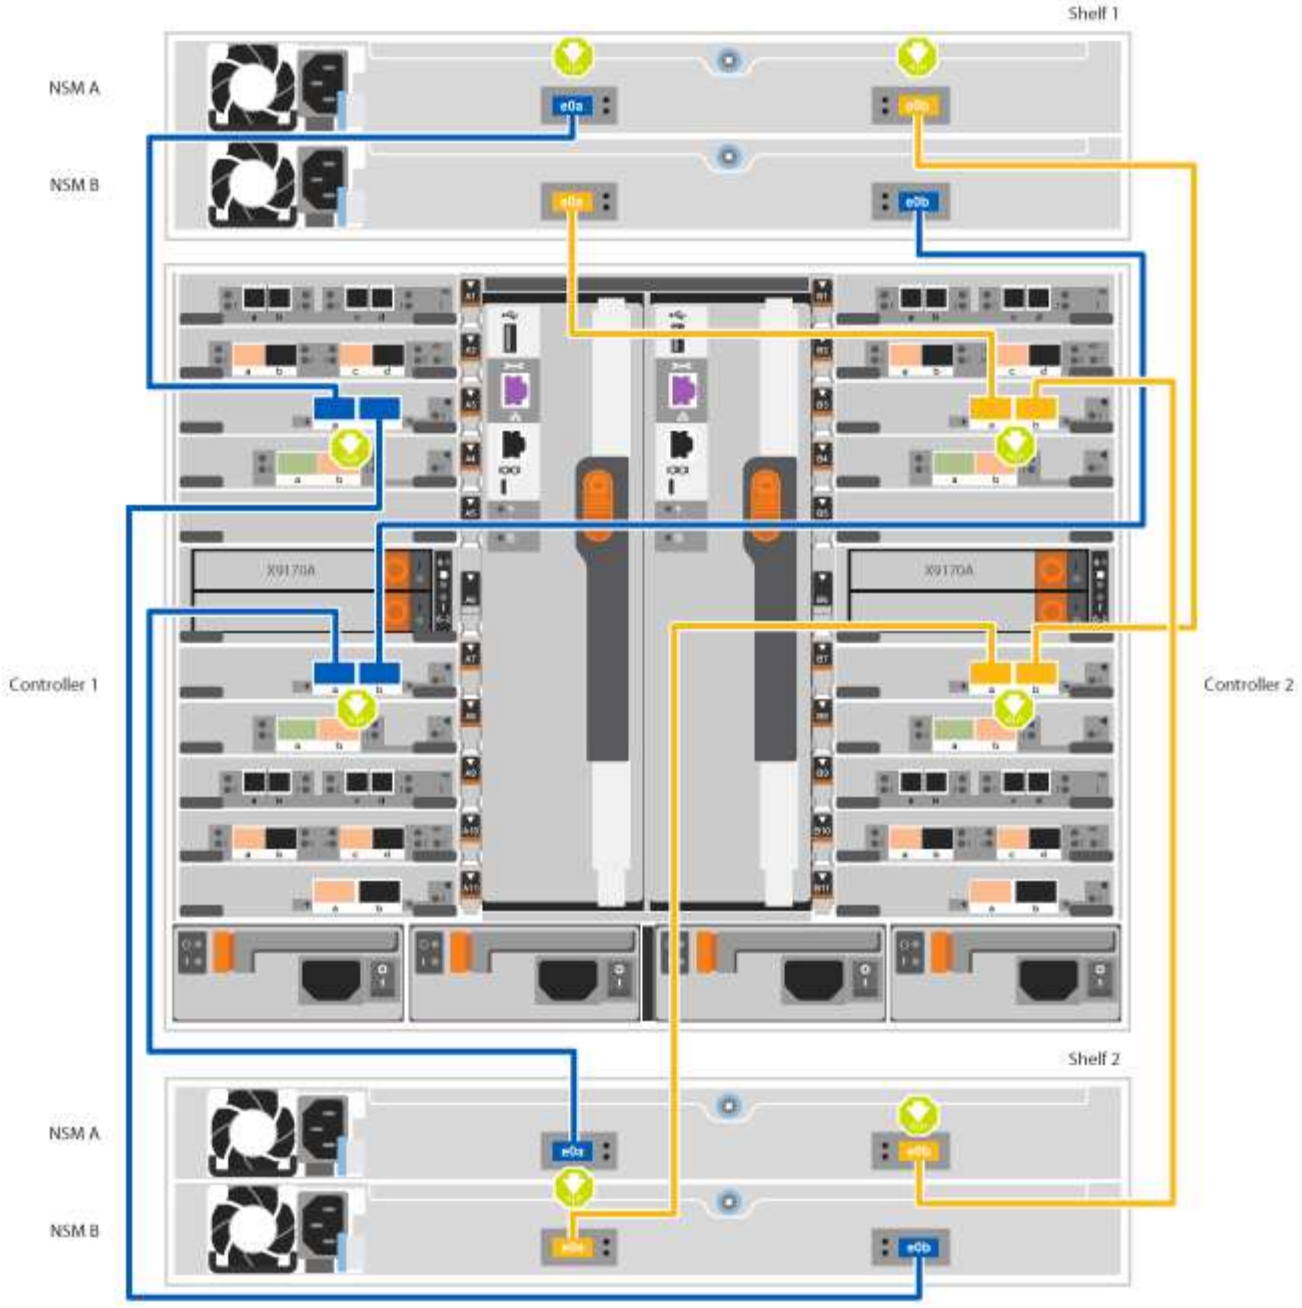

2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Step 5: Complete system setup and configuration

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: Completing system setup and configuration if network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Use the following animation to set one or more drive shelf IDs:

If your system has NS224 drive shelves, the shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located.

Animation - Set SAS or NVMe drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Turn on the power switches to both nodes.

   Animation - Turn on the power to the controllers

   Initial booting may take up to eight minutes.

4. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

5. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

6. Select an ONTAP icon listed to discover:

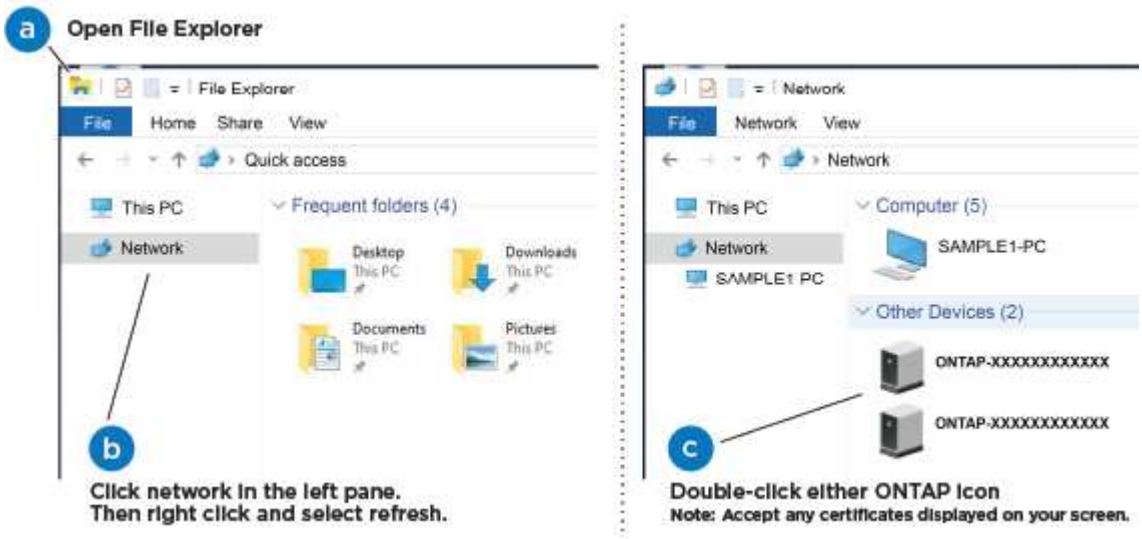

   a. Open File Explorer.
   b. Click **Network** in the left pane and right-click and select **refresh**.
   c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

      XXXXX is the system serial number for the target node.

   System Manager opens.

7. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

   ONTAP Configuration Guide

8. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

9. Verify the health of your system by running Config Advisor.
10. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet .

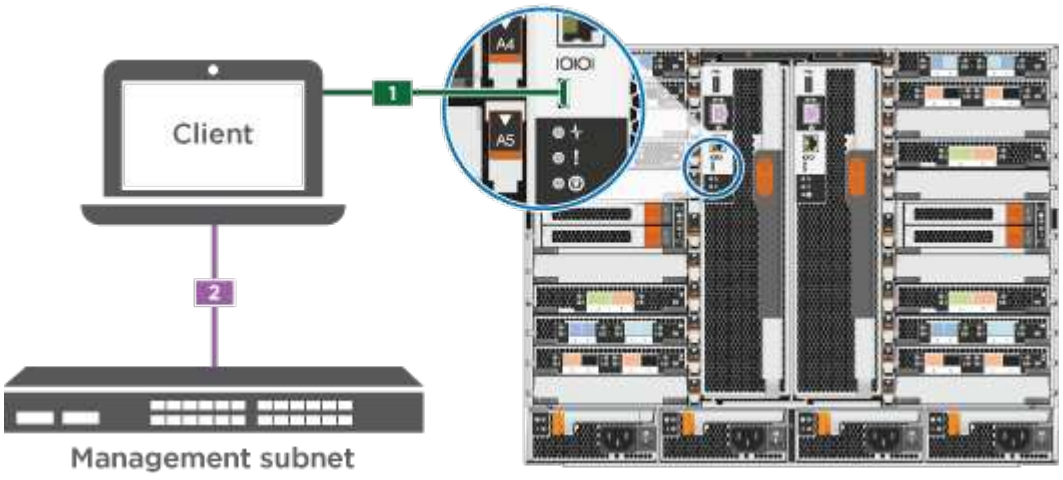

   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Use the following animation to set one or more drive shelf IDs:

   If your system has NS224 drive shelves, the shelves are pre-set to shelf ID 00 and 01. If you want to change the shelf IDs, you must create a tool to insert into the hole where button is located.

   Animation - Set SAS or NVMe drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
4. Turn on the power switches to both nodes.

   Animation - Turn on the power to the controllers

   Initial booting may take up to eight minutes.

5. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

6. Using System Manager on your laptop or console, configure your cluster:
   a. Point your browser to the node management IP address.

      The format for the address is https://x.x.x.x.

   b. Configure the system using the data you collected in the *NetApp ONTAP Configuration guide*.

      ONTAP Configuration Guide

7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain FAS9000 hardware

For the FAS9000 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Caching module

You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DCPM

The DCPM (destage controller power module) contains the NVRAM11 battery.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Fan

The fan cools the controller.

I/O module

The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.

LED USB

The LED USB module provides connectivity to console ports and system status.

NVRAM

The NVRAM module (Non-Volatile Random Access Memory) allows the controller to retain data across power cycles or system reboots.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real-time clock battery

A real time clock battery preserves system date and time information if the power is off.

X91148A module

The X91148A module is an I/O module that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.

Boot media

Overview of boot media replacement - FAS9000

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz`.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair does not require connection to a network to restore the `var` file system. The HA pair in a single chassis has an internal e0S connection, which is used to transfer `var` config between them.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct node:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy node* is the HA partner of the impaired node.

Check encryption key support and status

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

Step 1: Check if your version of ONTAP supports NetApp Volume Encryption

Check whether your ONTAP version supports NetApp Volume Encryption (NVE). This information is crucial for downloading the correct ONTAP image.

Steps

1. Determine if your ONTAP version supports encryption by running the following command:

   `version -v`

   If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Depending on whether NVE is supported on your system, take one of the following actions:
   - If NVE is supported, download the ONTAP image with NetApp Volume Encryption.
   - If NVE is not supported, download the ONTAP image **without** NetApp Volume Encryption.

Step 2: Determine if it is safe to shut down the controller

To safely shut down a controller, first identify whether the External Key Manager (EKM) or the Onboard Key Manager (OKM) is active. Then, verify the key manager in use, display the appropriate key information, and take action based on the status of the authentication keys.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manger is configured on your system, select one of the following options.

No key manager configured

You can safely shut down the impaired controller. Go to shutdown the impaired controller.

External or Onboard key manager configured

a. Enter the following query command to display the status of the authentication keys in your key manager.

```
security key-manager key query
```

b. Check the output for the value in the `Restored` column for your key manager.

This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Depending on whether your system is using the External Key Manager or Onboard Key Manager, select one of the following options.

External Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	You can safely shut down the impaired controller. Go to shutdown the impaired controller.
Anything other than `true`	a. Restore the external key management authentication keys to all nodes in the cluster using the following command: `security key-manager external restore` If the command fails, contact NetApp Support. b. Verify that the `Restored` column displays `true` for all authentication keys by entering the `security key-manager key query` command. If all the authentication keys are `true`, you can safely shut down the impaired controller. Go to shutdown the impaired controller.

Onboard Key Manager

Depending on the output value displayed in the `Restored` column, follow the appropriate steps.

Output value in `Restored` column	**Follow these steps…**
`true`	Manually back up the OKM information. a. Go to the advanced mode by entering `set -priv advanced` and then enter `Y` when prompted. b. Enter the following command to display the key management information: `security key-manager onboard show-backup` c. Copy the contents of the backup information to a separate file or your log file. You'll need it in disaster scenarios where you might need to manually recover OKM. d. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Output value in `Restored` column	Follow these steps…

Anything other than `true`	a. Enter the onboard security key-manager sync command:
	
	`security key-manager onboard sync`
	
	b. Enter the 32 character, alphanumeric onboard key management passphrase when prompted.
	
	If the passphrase cannot be provided, contact NetApp Support.
	
	c. Verify the `Restored` column displays `true` for all authentication keys:
	
	`security key-manager key query`
	
	d. Verify that the `Key Manager` type displays `onboard`, and then manually back up the OKM information.
	e. Enter the command to display the key management backup information:
	
	`security key-manager onboard show-backup`
	
	f. Copy the contents of the backup information to a separate file or your log file.
	
	You'll need it in disaster scenarios where you might need to manually recover OKM.
	
	g. You can safely shut down the impaired controller. Go to shutdown the impaired controller.

Shut down the impaired controller - FAS9000

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller displays…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

> This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired node.
NOTE: Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify –node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 3: Controller is in a two-node MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired node.

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Replace the boot media - FAS9000

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

Step 1: Remove the controller

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the orange button on the cam handle downward until it unlocks.

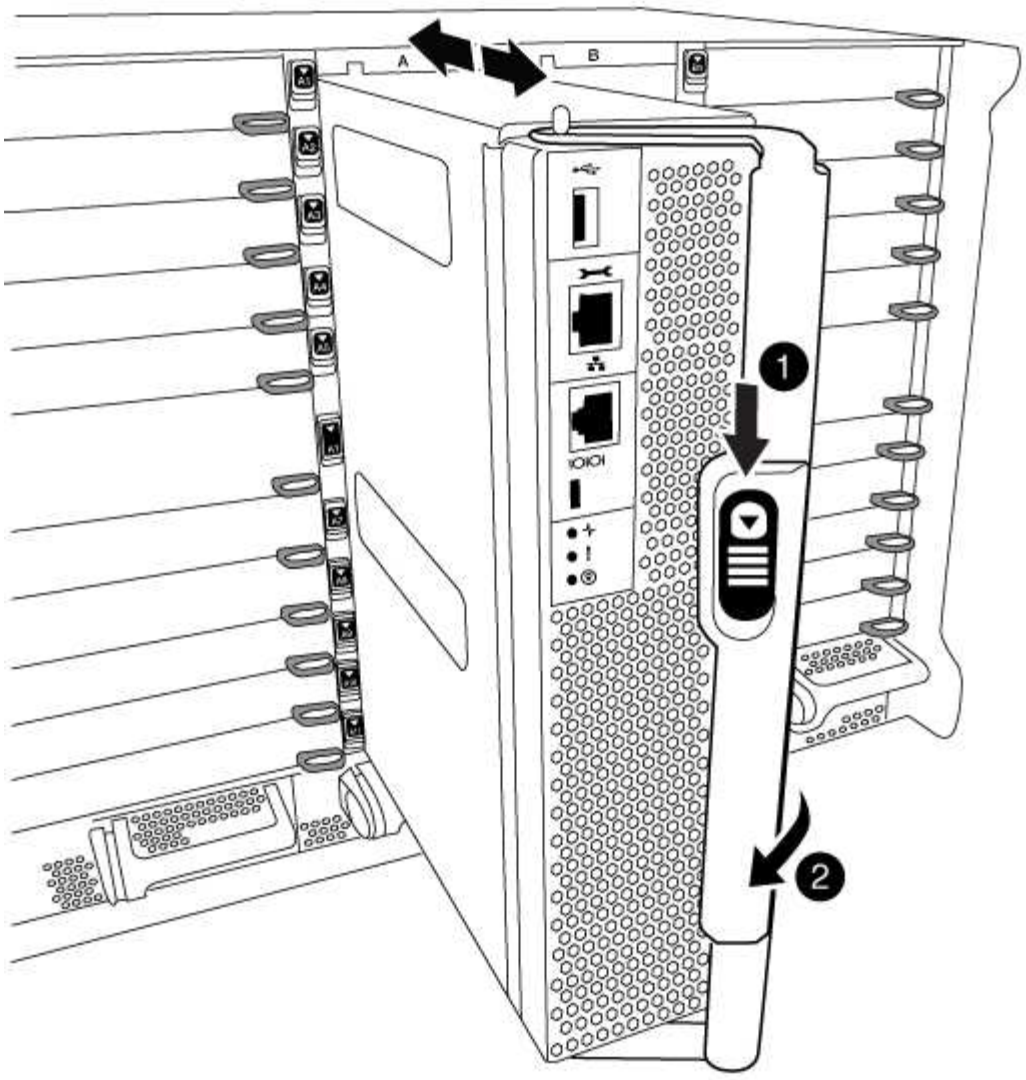

1	
Cam handle release button	
2	
Cam handle	

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

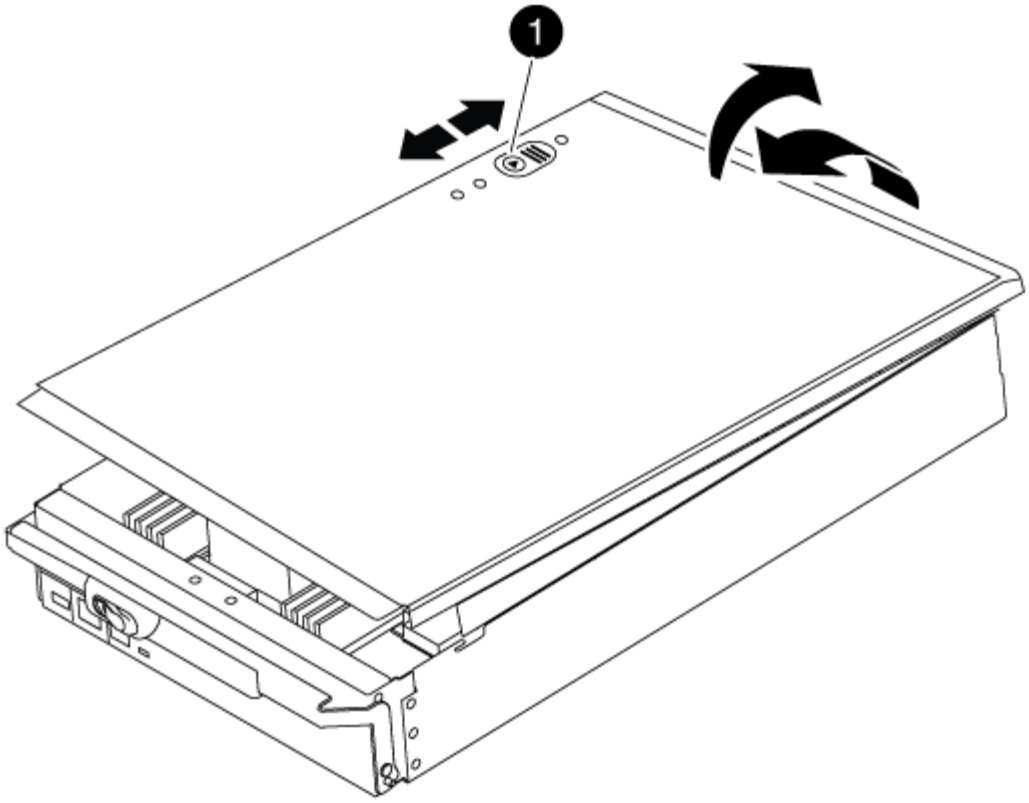

1	Controller module cover locking button

Step 2: Replace the boot media

Locate the boot media using the following illustration or the FRU map on the controller module:

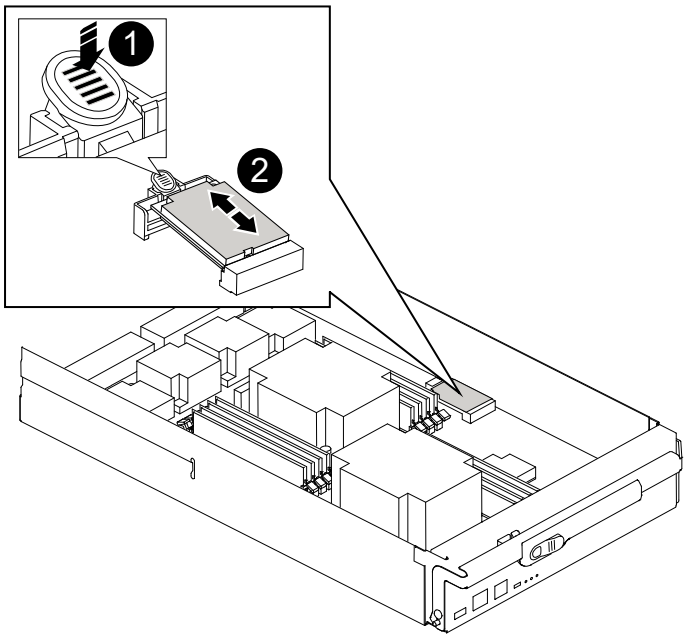

1	
Press release tab	
2	
Boot media	

1. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

2. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
3. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

4. Push the boot media down to engage the locking button on the boot media housing.
5. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the `var` file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the `var` file system.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Recable the controller module, as needed.
3. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB

console port.

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, and then push the cam handle to the closed position.

   The node begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then halt the node to boot to LOADER.

6. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

     The target port you configure is the target port you use to communicate with the impaired node from the healthy node during `var` file system restore with a network connection. You can also use the e0M port in this command.

   - If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
     - filer_addr is the IP address of the storage system.
     - netmask is the network mask of the management network that is connected to the HA partner.
     - gateway is the gateway for the network.
     - dns_addr is the IP address of a name server on your network.
     - dns_domain is the Domain Name System (DNS) domain name.

       If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

     Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

7. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:

   a. Boot to Maintenance mode: `boot_ontap maint`

   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`

   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Boot the recovery image - FAS9000

The procedure for booting the impaired node from the recovery image depends on whether the system is in a two-node MetroCluster configuration.

Option 1: Boot the recovery image in most systems

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

This procedure applies to systems that are not in a two-node MetroCluster configuration.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.

3. Restore the `var` file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy node to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the node to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the node.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

If your system has…	Then…

No network connection and is in a MetroCluster IP configuration	a. Press `n` when prompted to restore the backup configuration.
	b. Reboot the system when prompted by the system.
	c. Wait for the iSCSI storage connections to connect.
	You can proceed after you see the following messages:
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_auxiliary, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_partner, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_auxiliary, address: ip-address).`
	`date-and-time [node-name:iscsi.session.stateChanged:notice]: iSCSI session state is changed to Connected for the target iSCSI-target (type: dr_partner, address: ip-address).`
	d. Select the **Update flash from backup config** (sync flash) option from the displayed menu.
	If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the node to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:
   - If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed

- If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

*If you see…	Then…*
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner node. b. Confirm the target node is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner node.
8. Give back the node using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired node and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Option 2: Boot the recovery image in a two-node MetroCluster configuration

You must boot the ONTAP image from the USB drive and verify the environmental variables.

This procedure applies to systems in a two-node MetroCluster configuration.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. After the image is installed, start the restoration process:
   a. Press `n` when prompted to restore the backup configuration.
   b. Press `y` when prompted to reboot to start using the newly installed software.

      You should be prepared to interrupt the boot process when prompted.

4. As the system boots, press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message., and when the Boot Menu is displayed select option 6.
5. Verify that the environmental variables are set as expected.
   a. Take the node to the LOADER prompt.

b. Check the environment variable settings with the `printenv` command.

c. If an environment variable is not set as expected, modify it with the `setenv environment-variable-name changed-value` command.

d. Save your changes using the `savenv` command.

e. Reboot the node.

Switch back aggregates in a two-node MetroCluster configuration - FAS9000

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Restore encryption - FAS9000

Restore encryption on the replacement boot media.

You must complete steps specific to systems that have Onboard Key Manager (OKM), NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) enabled using the settings that you captured at the beginning of the boot media replace procedure.

Depending on which a key manger is configured on your system, select one of the following options to restore it from the boot menu.

- Option 1: Restore the Onboard Key Manager configuration
- Option 2: Restore the External Key Manager configuration

Option 1: Restore the Onboard Key Manager configuration

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

- Make sure you have following information while restoring the OKM configuration:
  - Cluster-wide passphrase entered while enabling onboard key management.
  - Backup information for the Onboard Key Manager.
- Perform the How to verify onboard key management backup and cluster-wide passphrase procedure before proceeding.

Steps

1. Connect the console cable to the target controller.
2. From the ONTAP boot menu select the appropriate option from the boot menu.

ONTAP version	Select this option
ONTAP 9.8 or later	Select option 10. **Show example boot menu** Please choose one of the following: (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management. Selection (1-11)? 10

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

3. Confirm that you want to continue the recovery process.

   Show example prompt

   ```
   This option must be used only in disaster recovery procedures. Are you
   sure? (y or n):
   ```

4. Enter the cluster-wide passphrase twice.

   While entering the passphrase the console will not show any input.

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information.

   a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line.

Show example prompt

```
Enter the backup data:

--------------------------BEGIN BACKUP--------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

---------------------------END BACKUP---------------------------
```

b. Press the enter key twice at the end of the input.

   The recovery process completes.

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
******************
* Select option "(1) Normal Boot." to complete recovery process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
******************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets.`
Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
***************
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
***************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message.

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

8. From the partner node, giveback the partner controller by entering the following command.

   `storage failover giveback -fromnode local -only-cfo-aggregates true.`

9. After booting with only the CFO aggregate, run the following command.

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted volumes
are not brought online automatically, they can be brought online
manually using the "volume online -vserver <vserver> -volume
<volume_name>" command.
```

If the sync is successful the cluster prompt is returned with no additional messages. If the sync fails an error message appears before returning to the cluster prompt. Do not continue until the the error is corrected and the sync runs successfully.

11. Ensure that all keys are synced by entering the following command.

    ```
    security key-manager key query -restored false.

    There are no entries matching your query.
    ```

    No results should appear when filtering for false in the restored parameter.

12. Giveback the node from the partner by entering the following command.

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback, if you disabled it, by entering the following command.

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation by entering the following command.

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Option 2: Restore the External Key Manager configuration

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

You need the following information for restoring the External Key Manager (EKM) configuration.

- A copy of the /cfcard/kmip/servers.cfg file from another cluster node or the following information:
  - The KMIP server address.
  - The KMIP port.
- A copy of the `/cfcard/kmip/certs/client.crt` file from another cluster node or the client certificate.

- A copy of the `/cfcard/kmip/certs/client.key` file from another cluster node or the client key.
- A copy of the `/cfcard/kmip/certs/CA.pem` file from another cluster node or the KMIP server CA(s).

Steps

1. Connect the console cable to the target controller.
2. Select option 11 from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. When prompted, confirm you have gathered the required information.

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file? {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. When prompted, enter the client and server information.

   Show prompt

   ```
   Enter the client certificate (client.crt) file contents:
   Enter the client key (client.key) file contents:
   Enter the KMIP server CA(s) (CA.pem) file contents:
   Enter the server configuration (servers.cfg) file contents:
   ```

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

After you enter the client and server information, the recovery process completes.

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****************************************************************
*******
* Select option "(1) Normal Boot." to complete the recovery process.
*
*****************************************************************
*******

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it.

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation by entering the following command.

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed part to NetApp - FAS9000

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the caching module or add/replace a core dump module - FAS9000

You must replace the caching module in the controller module when your system registers a single AutoSupport (ASUP) message that the module has gone offline; failure to do so results in performance degradation. If AutoSupport is not enabled, you can locate the failed caching module by the fault LED on the front of the module. You can also add or replace the 1TB, X9170A core dump module, which is required if you are installing NS224 drive shelves in an AFF A700 system.

Before you begin

- You must replace the failed component with a replacement FRU component you received from your provider.
- For instructions about hot swapping the caching module, see Hot-swapping a caching module.
- When removing, replacing, or adding caching or core dump modules, the target node must be halted to the LOADER.
- AFF A700 supports the 1TB core dump module, X9170A, which is required if you are adding NS224 drive shelves.
- The core dump modules can be installed in slots 6-1 and 6-2. The recommended best practice is to install the module in slot 6-1.
- The X9170A core dump module is not hot-swappable.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Replace or add a caching module

The NVMe SSD Flash Cache modules (Flash Cache or caching modules) are separate modules. They are located in the front of the NVRAM module. To replace or add a caching module, locate it on the rear of the system on slot 6, and then follow the specific sequence of steps to replace it.

Before you begin

Your storage system must meet certain criteria depending on your situation:

- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- The target node must be at the LOADER prompt before adding or replacing the caching module.
- The replacement caching module must have the same capacity as the failed caching module, but can be from a different supported vendor.
- All other components in the storage system must be functioning properly; if not, you must contact technical support.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed caching module, in slot 6, by the lit amber Attention LED on the front of the caching module.
3. Remove the caching module:

> If you are adding another caching module to your system, remove the blank module and go to the next step.

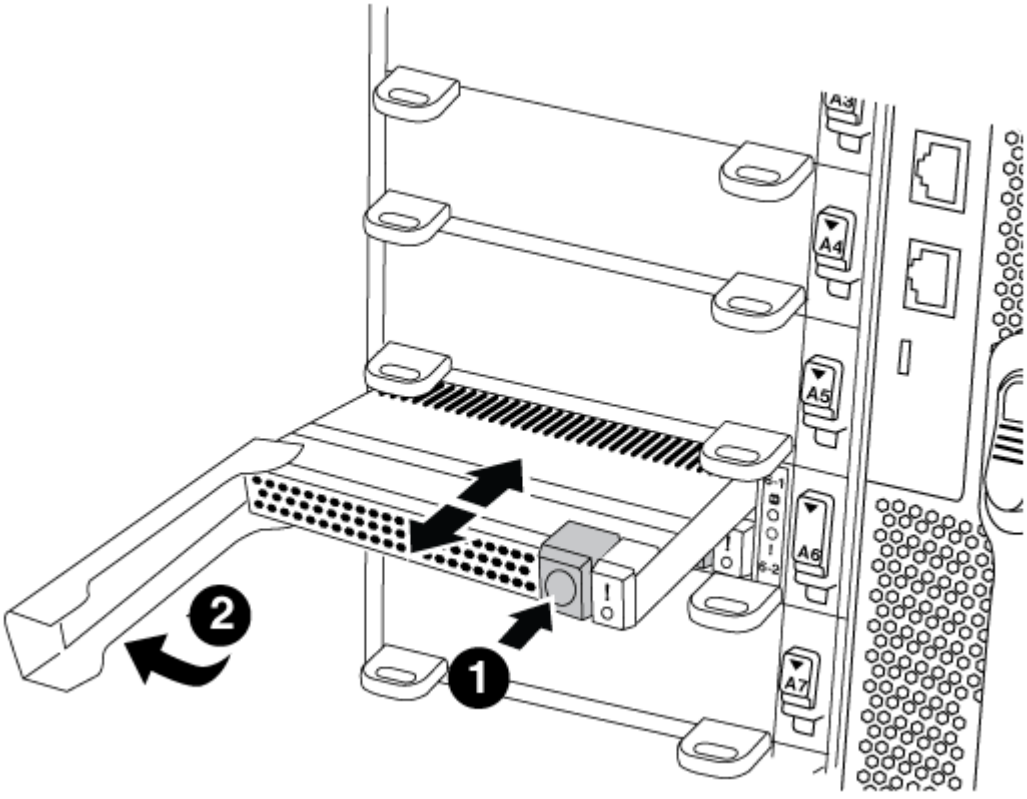

1	Orange release button.
2	Caching module cam handle.

a. Press the orange release button on the front of the caching module.

> Do not use the numbered and lettered I/O cam latch to eject the caching module. The numbered and lettered I/O cam latch ejects the entire NVRAM10 module and not the caching module.

b. Rotate the cam handle until the caching module begins to slide out of the NVRAM10 module.

c. Gently pull the cam handle straight toward you to remove the caching module from the NVRAM10 module.

   Be sure to support the caching module as you remove it from the NVRAM10 module.

4. Install the caching module:

   a. Align the edges of the caching module with the opening in the NVRAM10 module.

   b. Gently push the caching module into the bay until the cam handle engages.

   c. Rotate the cam handle until it locks into place.

Step 3: Add or replace an X9170A core dump module

The 1TB cache core dump, X9170A, is only used in the AFF A700 systems. The core dump module cannot be hot-swapped. The core dump module typically is located in the front of the NVRAM module in slot 6-1 in the rear of the system. To replace or add the core dump module, locate slot 6-1, and then follow the specific sequence of steps to add or replace it.

Before you begin

- Your system must be running ONTAP 9.8 or later in order to add a core dump module.
- The X9170A core dump module is not hot-swappable.
- The target node must be at the LOADER prompt before adding or replacing the code dump module.
- You must have received two X9170 core dump modules; one for each controller.
- All other components in the storage system must be functioning properly; if not, you must contact technical support.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you are replacing a failed core dump module, locate and remove it:

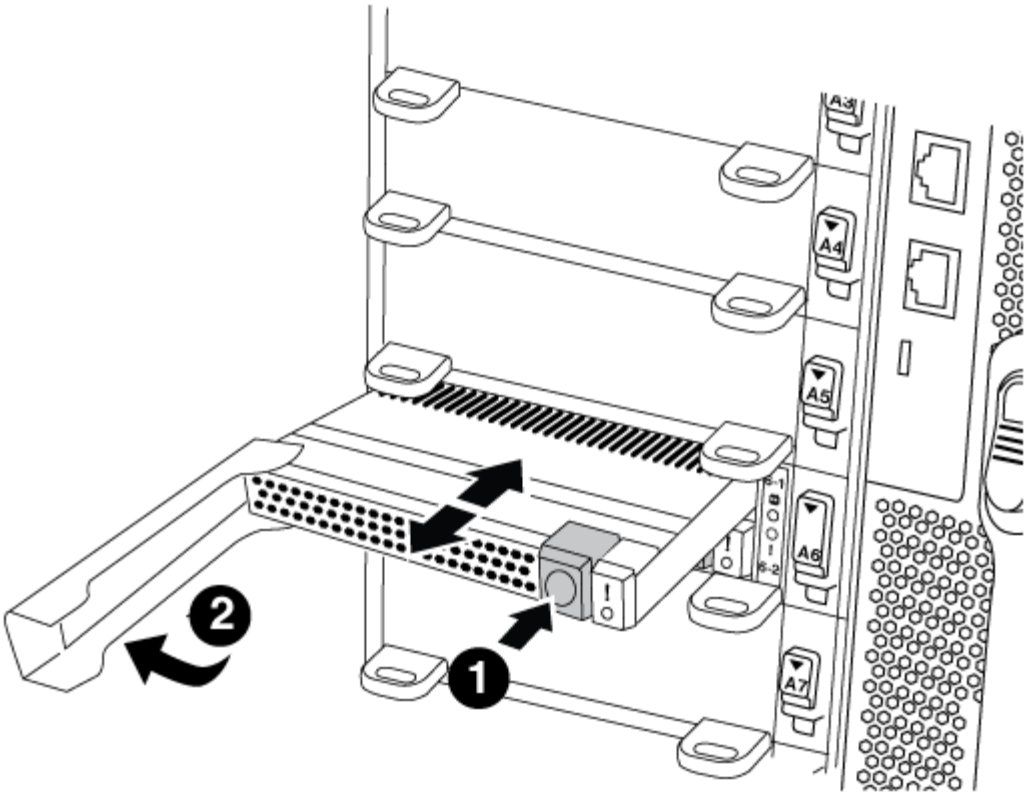

1	Orange release button.
2	Core dump module cam handle.

a. Locate the failed module by the amber Attention LED on the front of the module.

b. Press the orange release button on the front of the core dump module.

> Do not use the numbered and lettered I/O cam latch to eject the core dump module. The numbered and lettered I/O cam latch ejects the entire NVRAM10 module and not the core dump module.

c. Rotate the cam handle until the core dump module begins to slide out of the NVRAM10 module.

d. Gently pull the cam handle straight toward you to remove the core dump module from the NVRAM10 module and set it aside.

   Be sure to support the core dump module as you remove it from the NVRAM10 module.

3. Install the core dump module:

   a. If you are installing a new core dump module, remove the blank module from slot 6-1.

   b. Align the edges of the core dump module with the opening in the NVRAM10 module.

   c. Gently push the core dump module into the bay until the cam handle engages.

   d. Rotate the cam handle until it locks into place.

Step 4: Reboot the controller after FRU replacement

After you replace the FRU, you must reboot the controller module.

Step

1. To boot ONTAP from the LOADER prompt, enter `bye`.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    ------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a caching module - FAS9000

The NVMe SSD Flash Cache modules (Flash Cache or caching modules) are located in the front of the NVRAM10 module in Slot 6 of FAS9000 systems only. Beginning with ONTAP 9.4, you can hot-swap the caching module of the same capacity from the same or different supported vendor.

Before you begin

Your storage system must meet certain criteria depending on your situation:

- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- The replacement caching module must have the same capacity as the failed caching module, but can be from a different supported vendor.
- All other components in the storage system must be functioning properly; if not, you must contact technical support.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed caching module, in slot 6, by the lit amber Attention LED on the front of the caching module.
3. Prepare the caching module slot for replacement as follows:

   a. For ONTAP 9.7 and earlier:

      i. Record the caching module capacity, part number, and serial number on the target node: `system node run local sysconfig -av 6`

      ii. In admin privilege level, prepare the target NVMe slot for replacement, responding `y` when prompted whether to continue: `system controller slot module replace -node` *`node_name -slot slot_number`* The following command prepares slot 6-2 on node1 for replacement, and displays a message that it is safe to replace:

```
::> system controller slot module replace -node node1 -slot 6-2

Warning: NVMe module in slot 6-2 of the node node1 will be powered
off for replacement.
Do you want to continue? (y|n): `y`

The module has been successfully powered off.  It can now be
safely replaced.
After the replacement module is inserted, use the "system
controller slot module insert" command to place the module into
service.
```

iii. Display the slot status with the system controller slot module show command.

The NVMe slot status displays waiting-for-replacement in the screen output for the caching module that needs replacing.

b. For ONTAP 9.8 and later:

i. Record the caching module capacity, part number, and serial number on the target node: `system node run local sysconfig -av 6`

ii. In admin privilege level, prepare the target NVMe slot for removal, responding `y` when prompted whether to continue: `system controller slot module remove -node node_name -slot slot_number` The following command prepares slot 6-2 on node1 for removal, and displays a message that it is safe to remove:

```
::> system controller slot module remove -node node1 -slot 6-2

Warning: SSD module in slot 6-2 of the node node1 will be powered
off for removal.
Do you want to continue? (y|n): `y`

The module has been successfully removed from service and powered
off. It can now be safely removed.
```

iii. Display the slot status with the `system controller slot module show` command.

The NVMe slot status displays `powered-off` in the screen output for the caching module that needs replacing.

ⓘ See the Command man pages for your version of ONTAP for more details.

4. Remove the caching module:

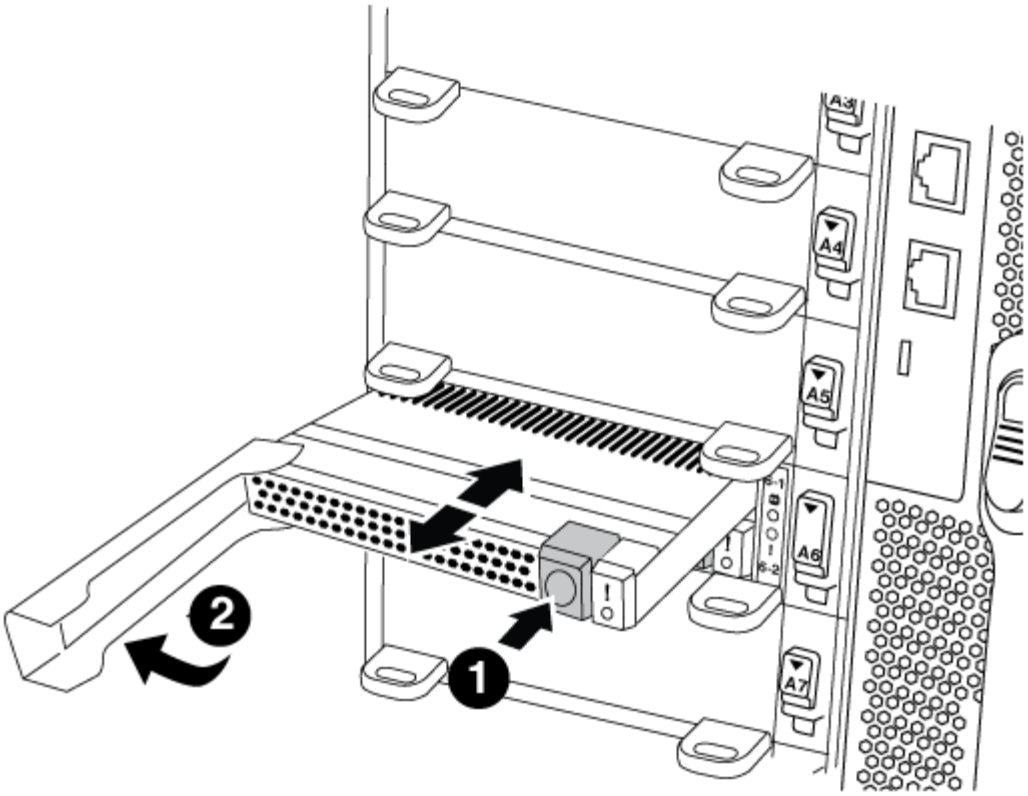

1	Orange release button.
2	Caching module cam handle.

a. Press the orange release button on the front of the caching module.

> Do not use the numbered and lettered I/O cam latch to eject the caching module. The numbered and lettered I/O cam latch ejects the entire NVRAM10 module and not the caching module.

b. Rotate the cam handle until the caching module begins to slide out of the NVRAM10 module.

c. Gently pull the cam handle straight toward you to remove the caching module from the NVRAM10 module.

   Be sure to support the caching module as you remove it from the NVRAM10 module.

5. Install the caching module:

   a. Align the edges of the caching module with the opening in the NVRAM10 module.

   b. Gently push the caching module into the bay until the cam handle engages.

   c. Rotate the cam handle until it locks into place.

6. Bring the replacement caching module online by using the `system controller slot module insert` command as follows:

   The following command prepares slot 6-2 on node1 for power-on, and displays a message that it is

powered on:

```
::> system controller slot module insert -node node1 -slot 6-2

Warning: NVMe module in slot 6-2 of the node localhost will be powered
on and initialized.
Do you want to continue? (y|n): `y`

The module has been successfully powered on, initialized and placed into
service.
```

7. Verify the slot status using the `system controller slot module show` command.

   Make sure that command output reports status for slot 6-1 or 6-2 as `powered-on` and ready for operation.

8. Verify that the replacement caching module is online and recognized, and then visually confirm that the amber attention LED is not lit: `sysconfig -av slot_number`

   If you replace the caching module with a caching module from a different vendor, the new vendor name is displayed in the command output.

9. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - FAS9000

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - FAS9000

To replace the chassis, you must shutdown the controllers.

Option 1: Shut down the controllers

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.

- BMC accessability for each controller.

- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace
   chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict -sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Shut down a node in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

   If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Move and replace hardware - FAS9000

Move the fans, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Remove the power supplies

Steps

Removing the power supplies when replacing a chassis involves turning off, disconnecting, and then removing the power supply from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:

a. Turn off the power switch on the power supply.

b. Open the power cable retainer, and then unplug the power cable from the power supply.

c. Unplug the power cable from the power source.

3. Press and hold the orange button on the power supply handle, and then pull the power supply out of the chassis.

When removing a power supply, always use two hands to support its weight.

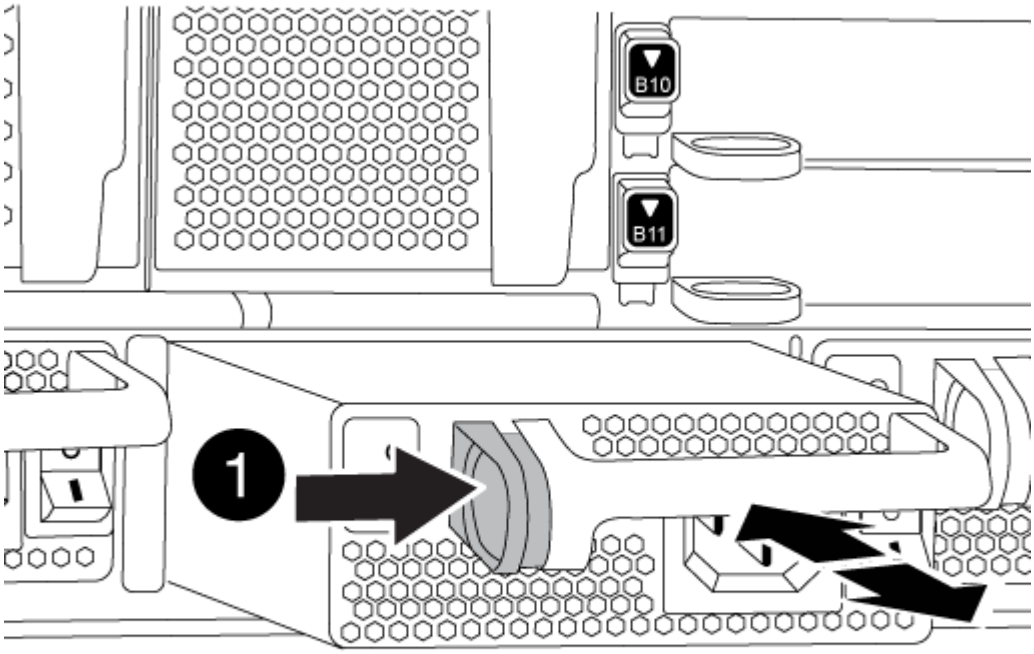

1	Locking button

4. Repeat the preceding steps for any remaining power supplies.

Step 2: Remove the fans

To remove the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

Steps

1. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
2. Press the orange button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

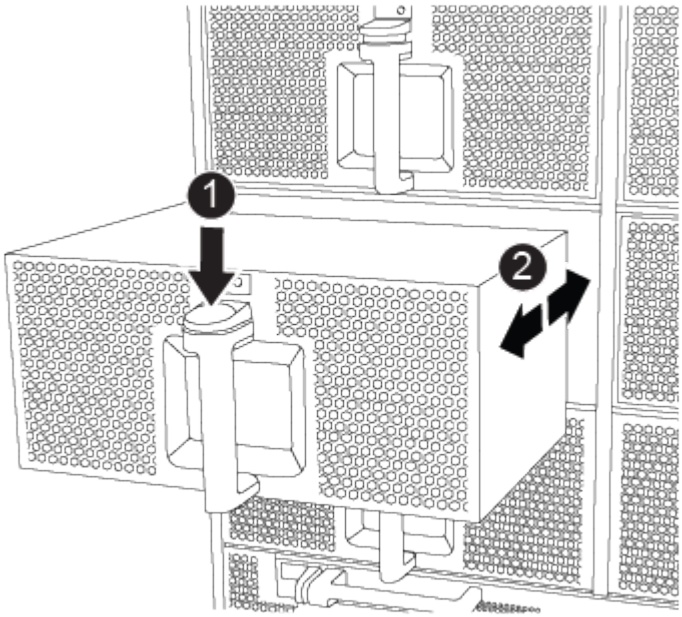

1	Orange release button

3. Set the fan module aside.
4. Repeat the preceding steps for any remaining fan modules.

Step 3: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the old chassis.

Steps

1. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
2. Slide the orange button on the cam handle downward until it unlocks.

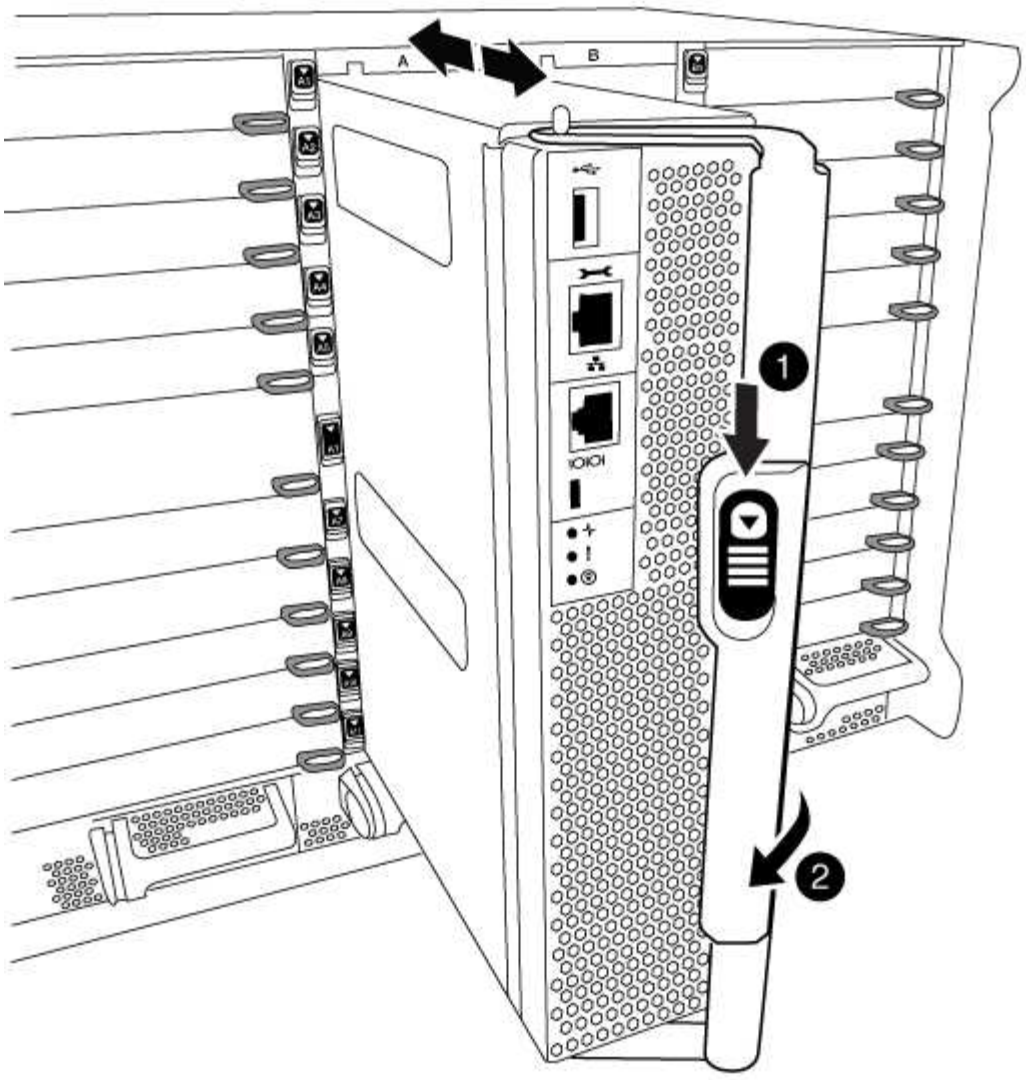

1	Cam handle release button
2	Cam handle

3. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

4. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 4: Remove the I/O modules

Steps

To remove I/O modules from the old chassis, including the NVRAM modules, follow the specific sequence of steps. You do not have to remove the Flash Cache module from the NVRAM module when moving it to a new chassis.

1. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

2. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

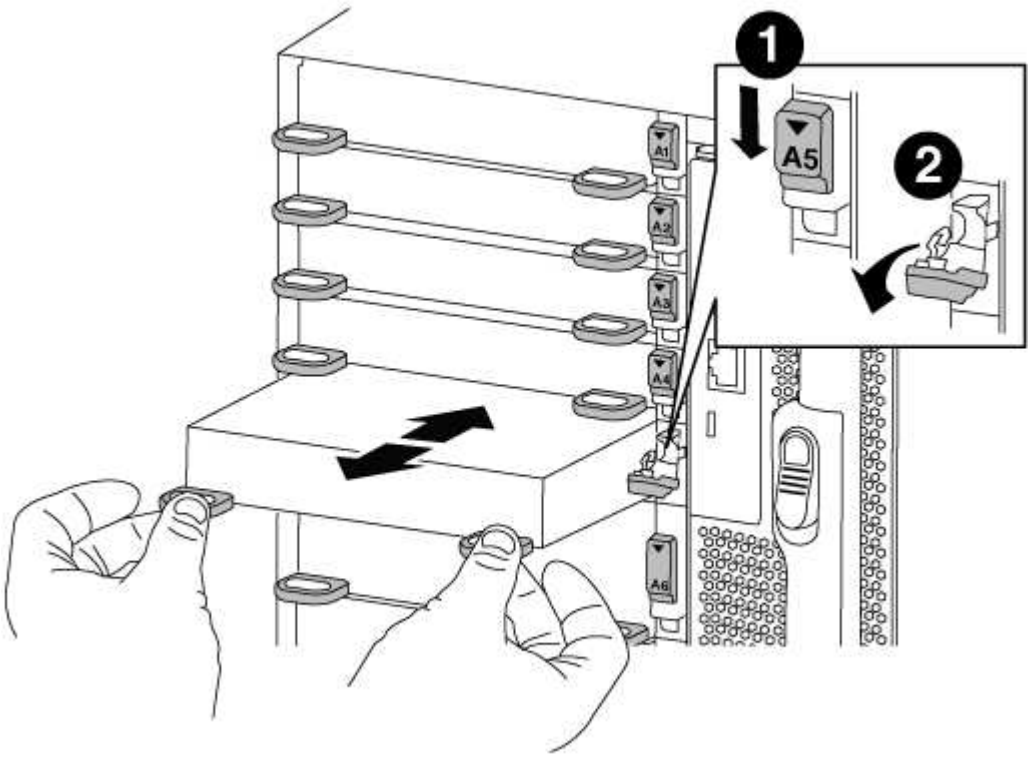

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

3. Set the I/O module aside.
4. Repeat the preceding step for the remaining I/O modules in the old chassis.

Step 5: Remove the De-stage Controller Power Module

Steps

You must remove the de-stage controller power modules from the old chassis in preparation for installing the replacement chassis.

1. Press the orange locking button on the module handle, and then slide the DCPM module out of the chassis.

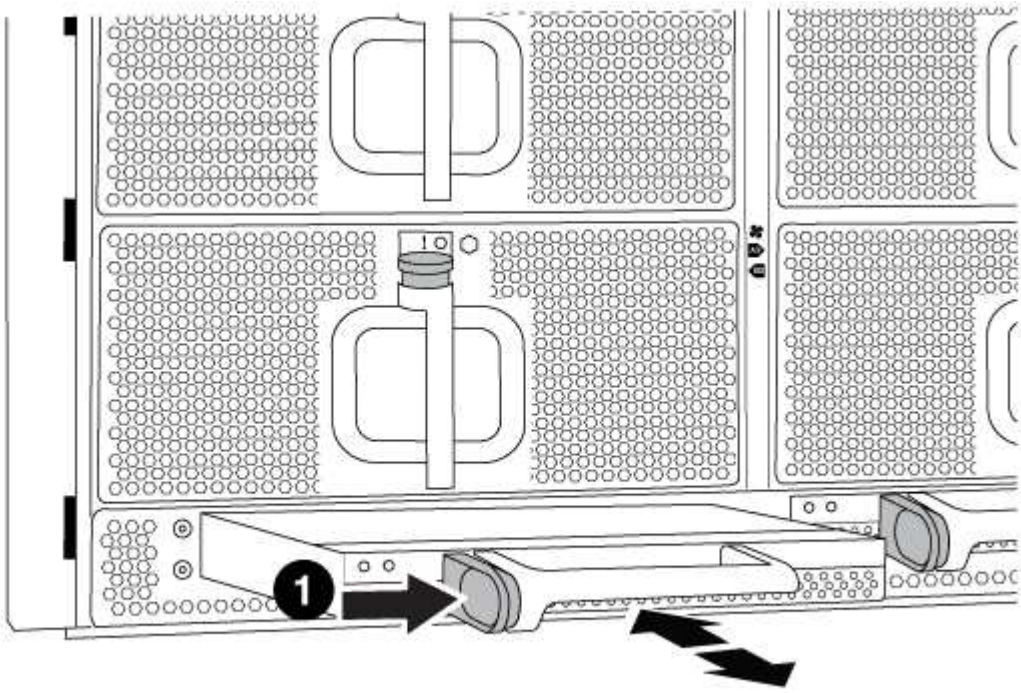

1	DCPM module orange locking button

2. Set the DCPM module aside in a safe place and repeat this step for the remaining DCPM module.

Step 6: Replace a chassis from within the equipment rack or system cabinet

Steps

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

1. Remove the screws from the chassis mount points.

   If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

2. With the help of two or three people, slide the old chassis off the rack rails in a system cabinet or *L* brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.

5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. Secure the rear of the chassis to the equipment rack or system cabinet.
8. If you are using the cable management brackets, remove them from the old chassis, and then install them on the replacement chassis.
9. If you have not already done so, install the bezel.

Step 7: Move the USB LED module to the new chassis

Steps

Once the new chassis is installed into the rack or cabinet, you must move the USB LED module from the old chassis to the new chassis.

1. Locate the USB LED module on the front of the old chassis, directly under the power supply bays.
2. Press the black locking button on the right side of the module to release the module from the chassis, and then slide it out of the old chassis.
3. Align the edges of the module with the USB LED bay at the bottom-front of the replacement chassis, and gently push the module all the way into the chassis until it clicks into place.

Step 8: Install the de-stage controller power module when replacing the chassis

Steps

Once the replacement chassis is installed into the rack or system cabinet, you must reinstall the de-stage controller power modules into it.

1. Align the end of the DCPM module with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

2. Repeat this step for the remaining DCPM module.

Step 9: Install fans into the chassis

Steps

To install the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

1. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

2. Repeat these steps for the remaining fan modules.
3. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Step 10: Install I/O modules

Steps

To install I/O modules, including the NVRAM/Flash Cache modules from the old chassis, follow the specific sequence of steps.

You must have the chassis installed so that you can install the I/O modules into the corresponding slots in the new chassis.

1. After the replacement chassis is installed in the rack or cabinet, install the I/O modules into their corresponding slots in the replacement chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage, and then push the I/O cam latch all the way up to lock the module in place.
2. Recable the I/O module, as needed.
3. Repeat the preceding step for the remaining I/O modules that you set aside.

   If the old chassis has blank I/O panels, move them to the replacement chassis at this time.

Step 11: Install the power supplies

Steps

Installing the power supplies when replacing a chassis involves installing the power supplies into the replacement chassis, and connecting to the power source.

1. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

2. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

   Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

3. Repeat the preceding steps for any remaining power supplies.

Step 12: Install the controller

Steps

After you install the controller module and any other components into the new chassis, boot it.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.

3. Connect the power supplies to different power sources, and then turn them on.
4. With the cam handle in the open position, slide the controller module into the chassis and firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle until it clicks into the locked position.

   Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.

   The controller module begins to boot as soon as it is fully seated in the chassis.

5. Repeat the preceding steps to install the second controller into the new chassis.
6. Boot each node to Maintenance mode:
   a. As each node starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

      If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

   b. From the boot menu, select the option for Maintenance mode.

Complete the restoration and replacement process - FAS9000

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

      The value for `HA-state` can be one of the following:

      - `ha`
      - `mcc`
      - `mcc-2n`
      - `mccip`
      - `non-ha`

b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.
4. Exit Maintenance mode: `halt`

   The LOADER prompt appears.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    ------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - FAS9000

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system has a V_StorageAttach license, you must refer to the additional required steps before performing this procedure.
- If your system is in an HA pair, the healthy node must be able to take over the node that is being replaced (referred to in this procedure as the “impaired node”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a node in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired node to the *replacement* node so that the *replacement* node will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* node is the node that is being replaced.
  - The *replacement* node is the new node that is replacing the impaired node.
  - The *healthy* node is the surviving node.

- You must always capture the node's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the controller module hardware - FAS9000

To replace the controller module hardware, you must remove the impaired node, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the orange button on the cam handle downward until it unlocks.

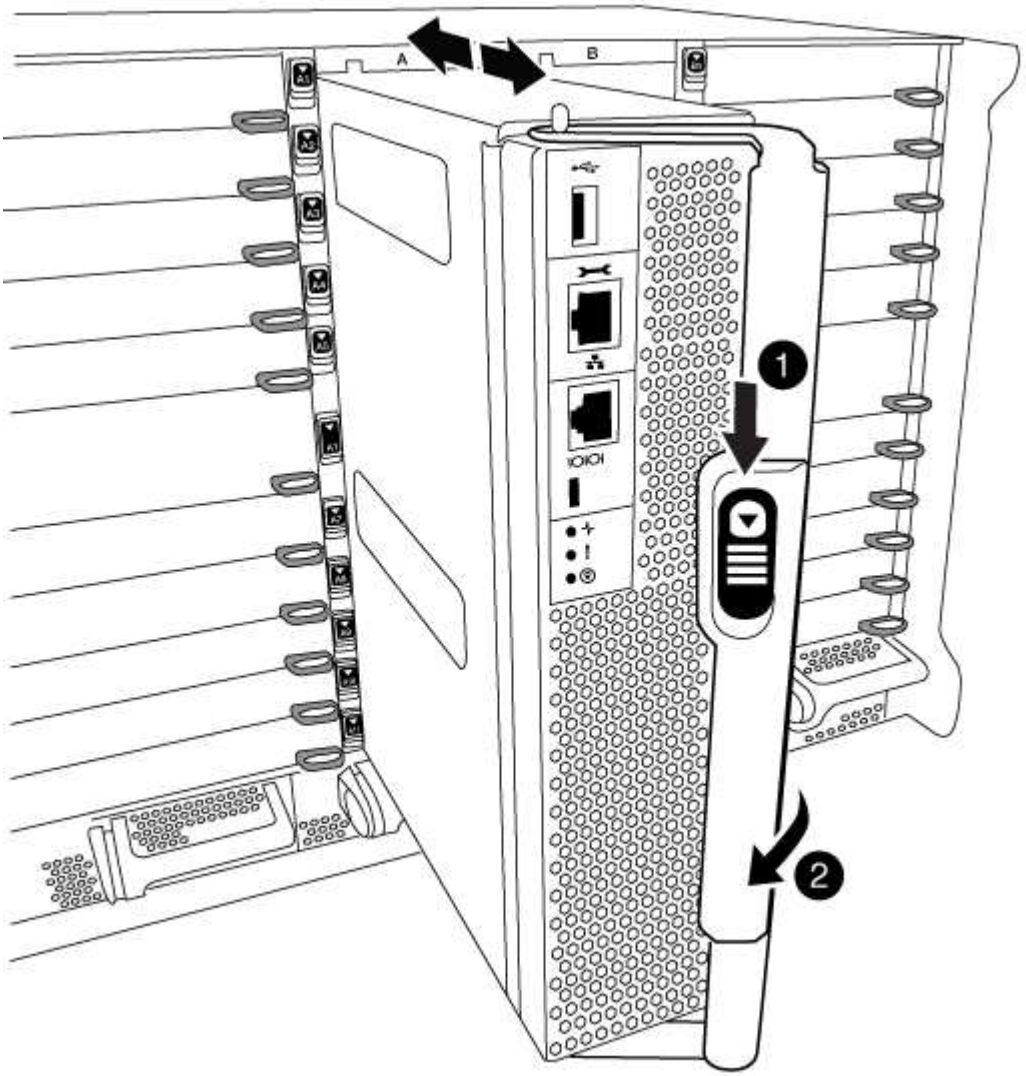

1	Cam handle release button
2	Cam handle

1. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

2. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

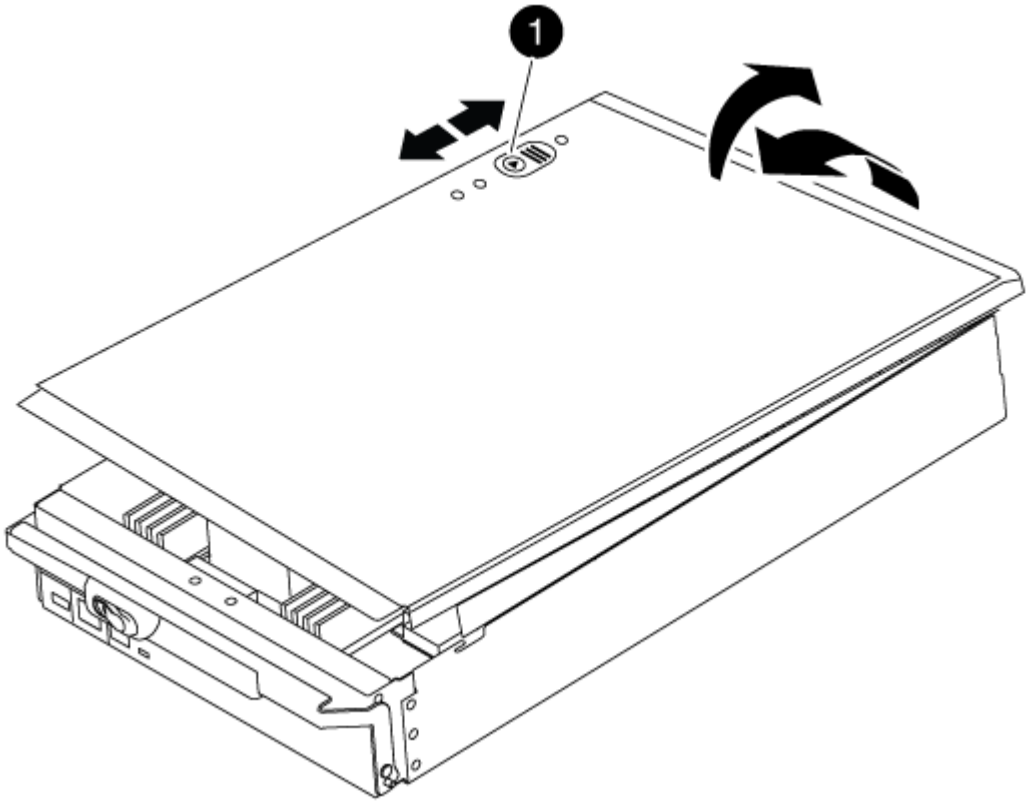

1	Controller module cover locking button

Step 2: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller and insert it in the new controller.

Steps

1. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

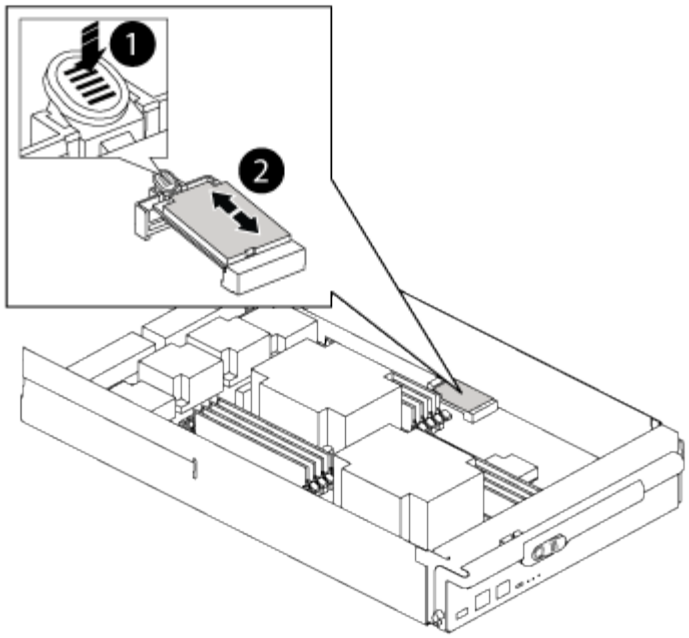

1	Press release tab
2	Boot media

2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.

Step 3: Move the system DIMMs

To move the DIMMs, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM,

and then slide the DIMM out of the slot.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

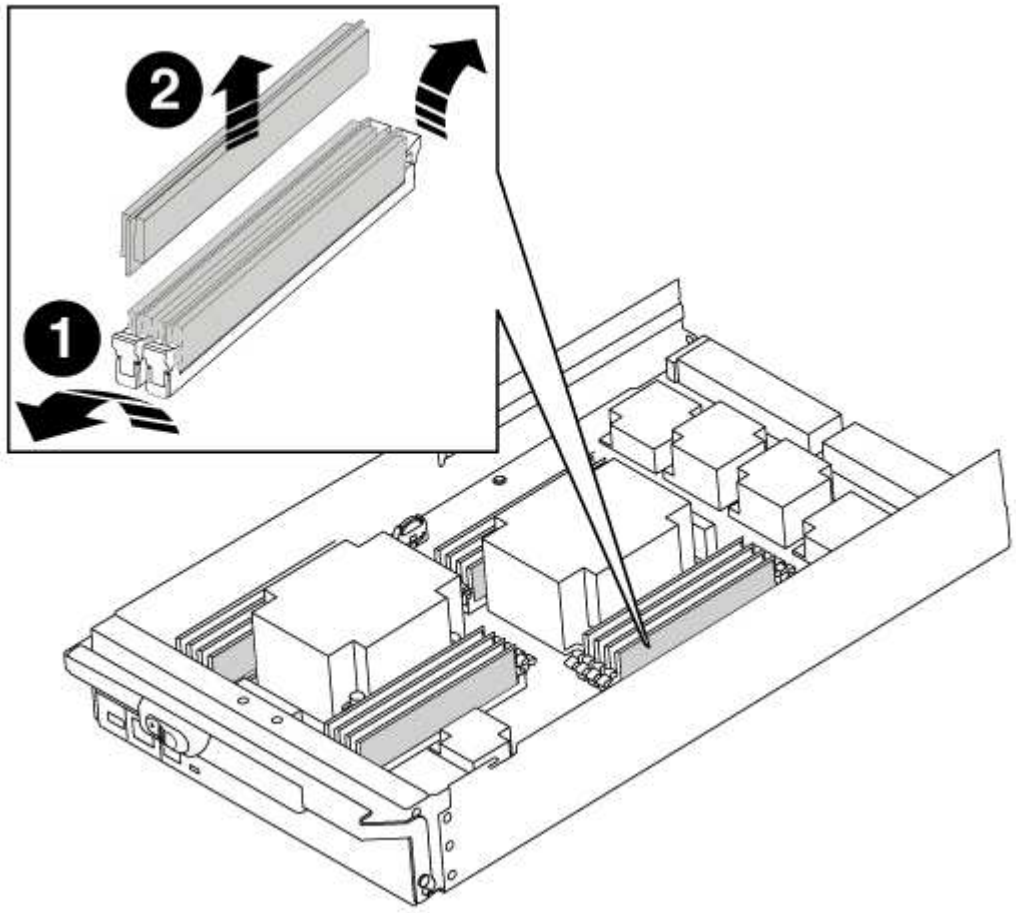

1	DIMM ejector tabs
2	DIMM

5. Locate the slot where you are installing the DIMM.
6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

8. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
9. Repeat these steps for the remaining DIMMs.

Step 4: Install the controller

After you install the components into the controller module, you must install the controller module back into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   d. Interrupt the boot process by pressing `Ctrl-C` when you see `Press Ctrl-C for Boot Menu`.
   e. Select the option to boot to Maintenance mode from the displayed menu.

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the controller's HA state

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`

- `mccip`
- `non-ha`

  a. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - FAS9000

Continue the replacement procedure by recabling the storage and confirming disk reassignment.

Step 1: Recable the system

Verify the controller module's storage and network connections.

Steps

1. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* node and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* node is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* node, boot the node, entering `y` if you are prompted to override the system ID due to a system ID mismatch.`boot_ontap`
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* node console and then, from the healthy node, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired node, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node                Partner         Possible     State Description
------------      ------------      --------
-------------------------------------
node1               node2           false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2               node1           -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy node, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the node:
   a. From the healthy node, give back the replaced node's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* node takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration Guide for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed

on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* node should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the node: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each node will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the node is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a node on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The *replacement* node is the current owner of the disks on the disaster site.

     Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each node is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each node: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy node: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - FAS9000

To complete the replacement procedure and restore your system to full operation, you must recable the storage, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Install licenses for the replacement node in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

Before you begin

If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

About this task

- Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

  Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

- The licenses keys must be in the 28-character format.
- You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.
- If the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verifying LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: (MetroCluster only): Switching back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a de-stage controller power module (DCPM) - FAS9000

To hot-swap a de-stage controller power module (DCPM), which contains the NVRAM10 battery, you must locate the failed DCPM module, remove it from the chassis, and install the replacement DCPM module.

You must have a replacement DCPM module in-hand before removing the failed module from the chassis and it must be replaced within five minutes of removal. Once the DCPM module is removed from the chassis, there is no shutdown protection for the controller module that owns the DCPM module, other than failover to the other controller module.

Step 1: Replace the DCPM module

To replace the DCPM module in your system, you must remove the failed DCPM module from the system and then replace it with a new DCPM module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel on the front of the system and set it aside.
3. Locate the failed DCPM module in the front of the system by looking for the Attention LED on the module.

   The LED will be steady amber if the module is faulty.

   > The DCPM module must be replaced in the chassis within five minutes of removal or the associated controller will shut down.

4. Press the orange locking button on the module handle, and then slide the DCPM module out of the chassis.

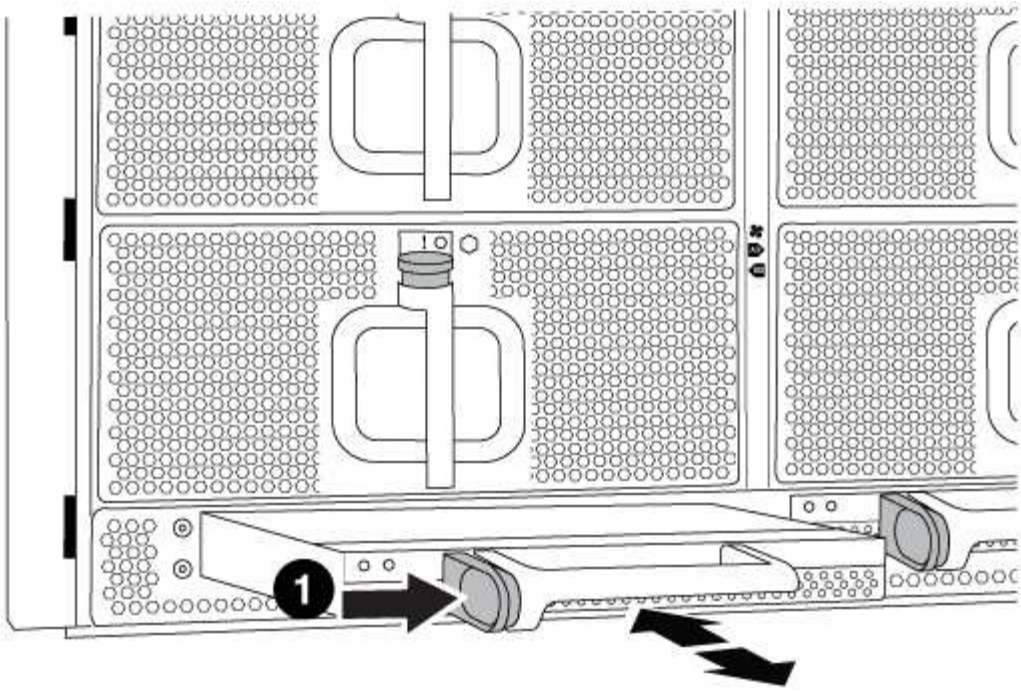

1	DCPM module orange locking button

5. Align the end of the DCPM module with the chassis opening, and then gently slide it into the chassis until it clicks into place.

> The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

The DCPM module LED lights when the module is fully seated into the chassis.

Step 2: Dispose of batteries

You must dispose of batteries according to the local regulations regarding battery recycling or disposal. If you cannot properly dispose of batteries, you must return the batteries to NetApp, as described in the RMA instructions that are shipped with the kit.

https://library.netapp.com/ecm/ecm_download_file/ECMP12475945

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS9000

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM

failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the orange button on the cam handle downward until it unlocks.

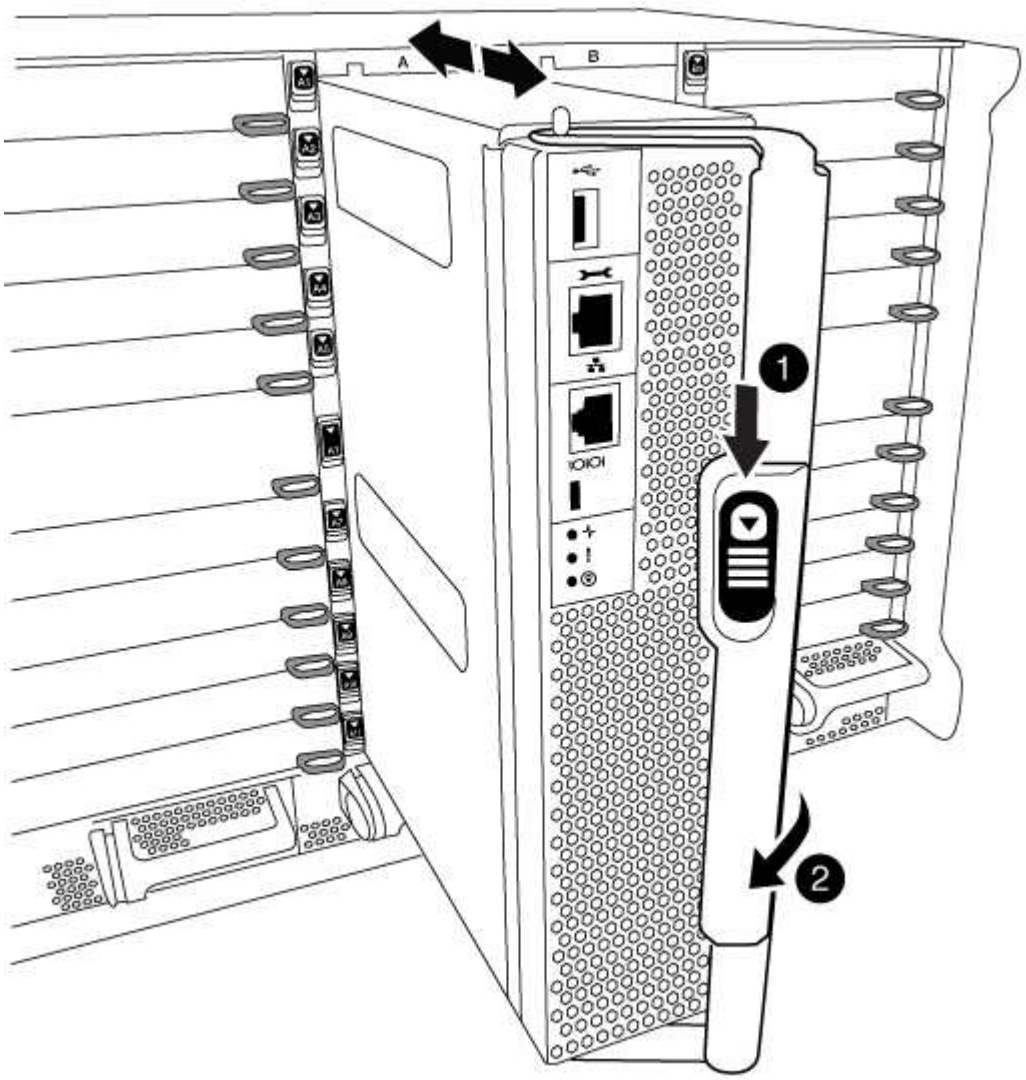

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

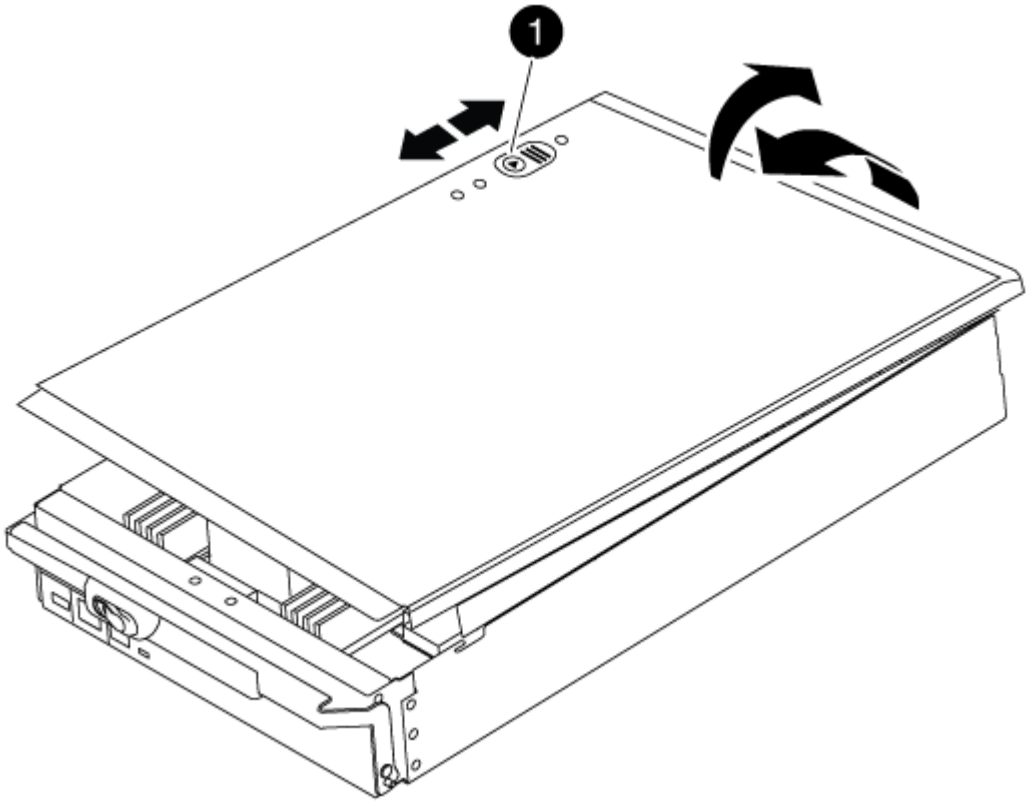

1	Controller module cover locking button

Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.

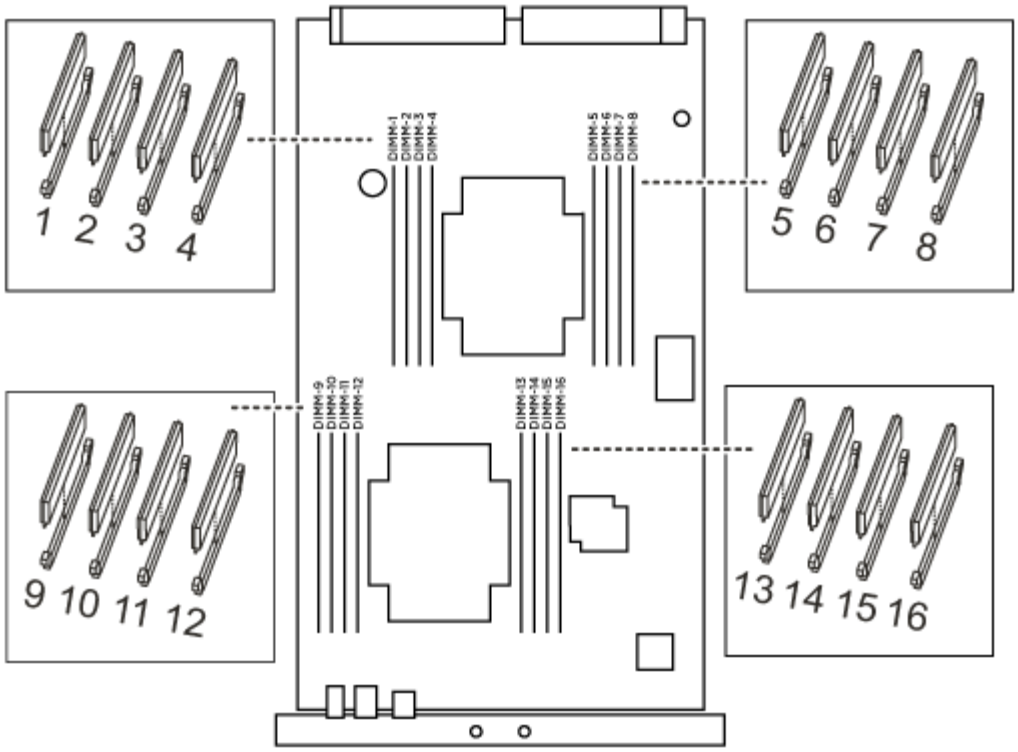

1. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

> Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

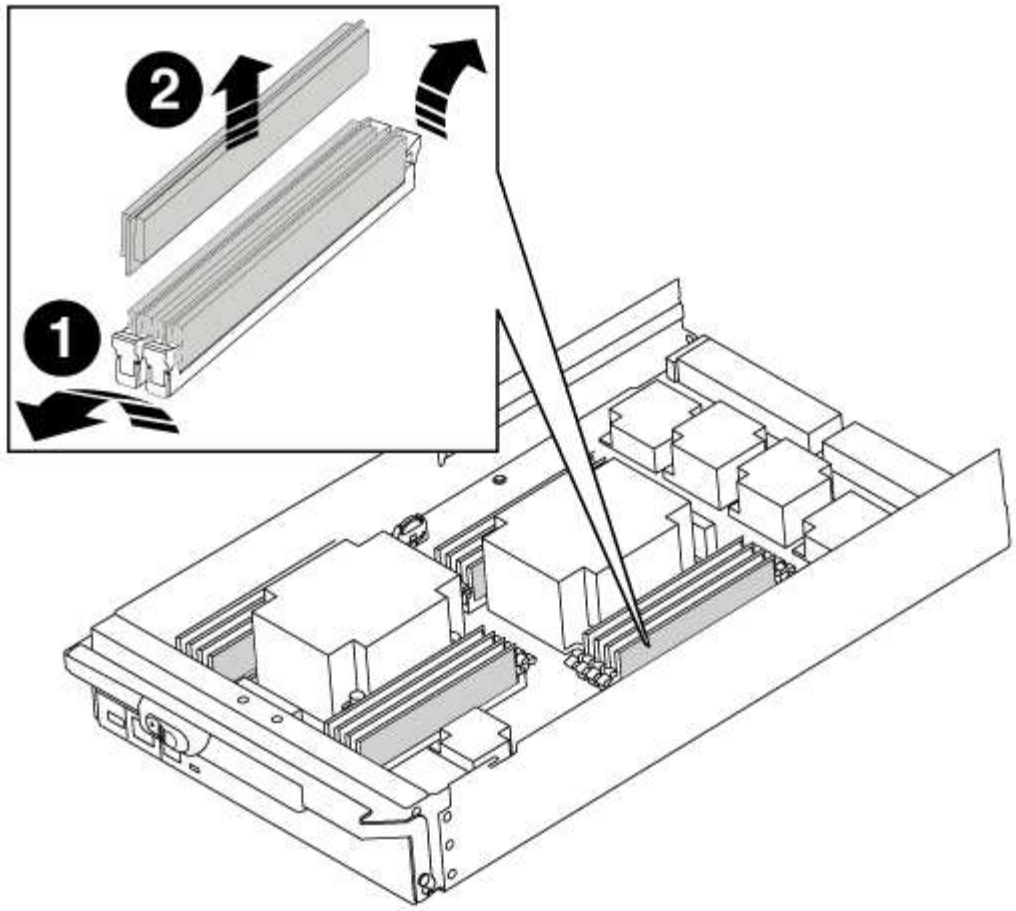

1	DIMM ejector tabs
2	DIMM

2. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

3. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

4. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
5. Close the controller module cover.

Step 4: Install the controller

After you install the components into the controller module, you must install the controller module back into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

4. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

5. Complete the reinstallation of the controller module:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a fan - FAS9000

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press the orange button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

> The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

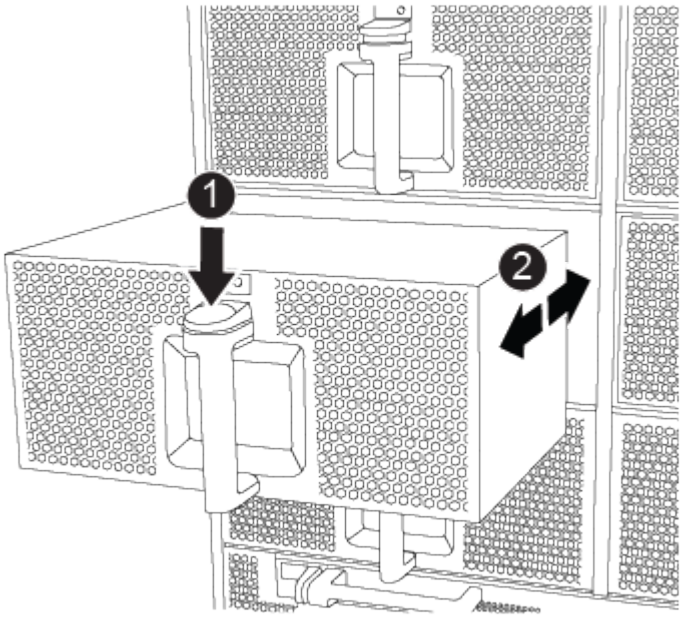

1	Orange release button

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the

chassis until it snaps into place.

When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an I/O module - FAS9000

To replace an I/O module, you must perform a specific sequence of tasks.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Replace I/O modules

To replace an I/O module, locate it within the chassis and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

3. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

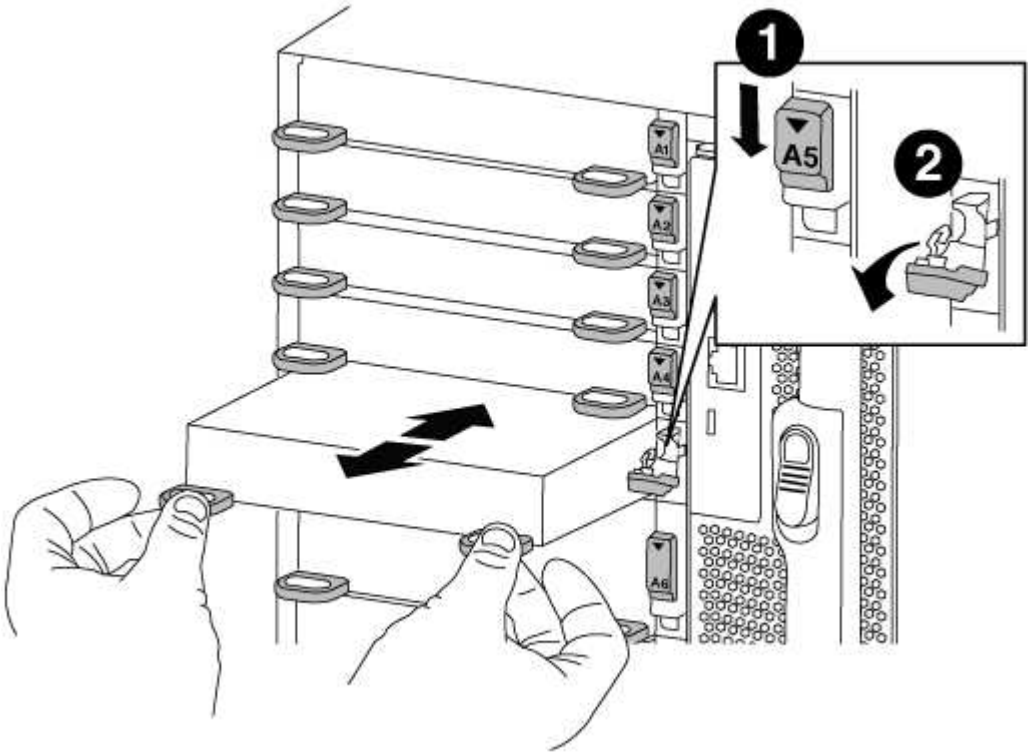

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Install the replacement I/O module into the chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.
6. Recable the I/O module, as needed.

Step 3: Reboot the controller after I/O module replacement

After you replace an I/O module, you must reboot the controller module.

> If the new I/O module is not the same model as the failed module, you must first reboot the BMC.

Steps

1. Reboot the BMC if the replacement module is not the same model as the old module:
   a. From the LOADER prompt, change to advanced privilege mode: `priv set advanced`
   b. Reboot the BMC: `sp reboot`
2. From the LOADER prompt, reboot the node: `bye`

   > This reinitializes the PCIe cards and other components and reboots the node.
3. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs or onboard ports, convert these ports to 10 GbE connections by using the `nicadmin convert` command from Maintenance mode.

   > Be sure to exit Maintenance mode after completing the conversion.
4. Return the node to normal operation:
   `storage failover giveback -ofnode` *`impaired_node_name`*
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

   > If your system is in a two-node MetroCluster configuration, you must switch back the aggregates as described in the next step.

Step 4: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an LED USB module - FAS9000

You can replace an LED USB module without interrupting service.

The FAS9000 or AFF A700 LED USB module provides connectivity to console ports and system status. Replacement of this module does not require tools.

Steps

1. Remove the old LED USB module:

   a. With the bezel removed, locate the LED USB module at the front of the chassis, on the bottom left side.
   b. Slide the latch to partially eject the module.
   c. Pull the module out of the bay to disconnect it from the midplane. Do not leave the slot empty.

2. Install the new LED USB module:

   a. Align the module to the bay with the notch in the corner of the module positioned near the slider latch on the chassis. The bay will prevent you from installing the module upside down.

b. Push the module into the bay until it is fully seated flush with the chassis.

   There is an audible click when the module is secure and connected to the midplane.

Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVRAM module or NVRAM DIMMs - FAS9000

The NVRAM module consists of the NVRAM10 and DIMMs and up to two NVMe SSD Flash Cache modules (Flash Cache or caching modules) per NVRAM module. You can replace a failed NVRAM module or the DIMMs inside the NVRAM module.

To replace a failed NVRAM module, you must remove it from the chassis, remove the Flash Cache module or modules from the NVRAM module, move the DIMMs to the replacement module, reinstall the Flash Cache module or modules, and install the replacement NVRAM module into the chassis.

Because the system ID is derived from the NVRAM module, if replacing the module, disks belonging to the system are reassigned to the new system ID.

Before you begin

- All disk shelves must be working properly.
- If your system is in an HA pair, the partner node must be able to take over the node associated with the NVRAM module that is being replaced.
- This procedure uses the following terminology:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy* node is the HA partner of the impaired node.
- This procedure includes steps for automatically or manually reassigning disks to the controller module associated with the new NVRAM module. You must reassign the disks when directed to in the procedure. Completing the disk reassignment before giveback can cause issues.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You cannot change any disks or disk shelves as part of this procedure.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a Two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 6 in the chassis and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Move the Flash Cache module from the old NVRAM module to the new NVRAM module:

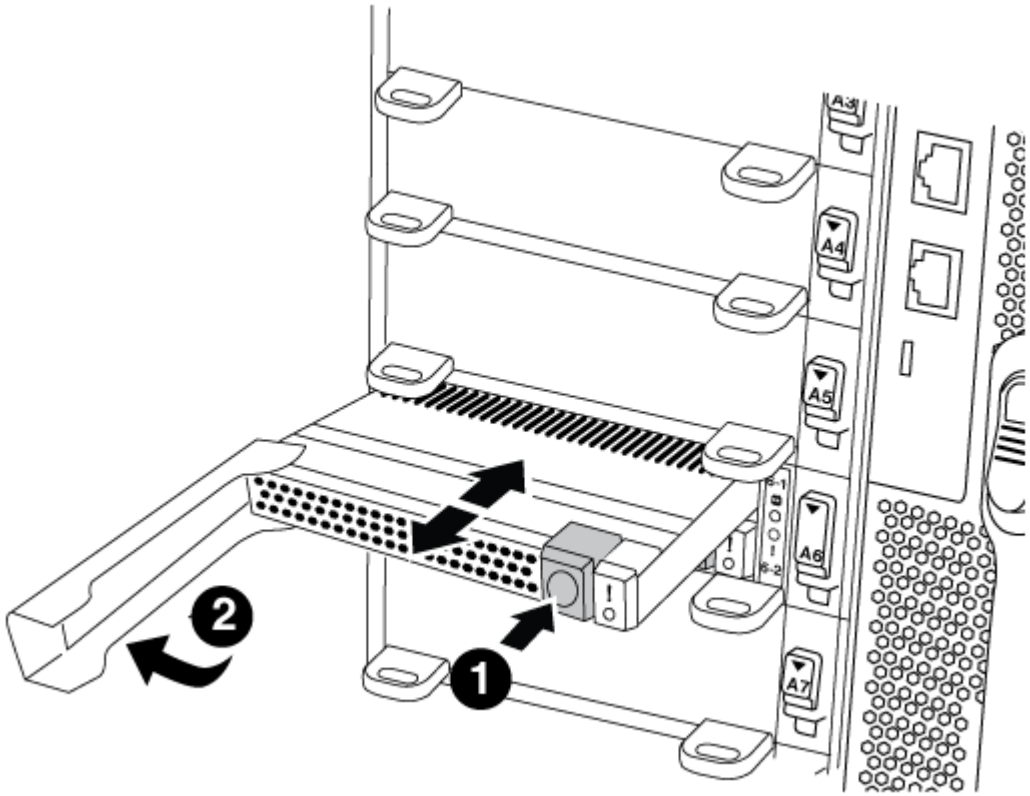

1	Orange release button (gray on empty Flash Cache modules)
2	Flash Cache cam handle

   a. Press the orange button on the front of the Flash Cache module.

      The release button on empty Flash Cache modules is gray.

   b. Swing the cam handle out until the module begins to slide out of the old NVRAM module.

   c. Grasp the module cam handle and slide it out of the NVRAM module and insert it into the front of the new NVRAM module.

   d. Gently push the Flash Cache module all the way into the NVRAM module, and then swing the cam handle closed until it locks the module in place.

3. Remove the target NVRAM module from the chassis:

   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

b. Rotate the cam latch down until it is in a horizontal position.

   The NVRAM module disengages from the chassis and moves out a few inches.

c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

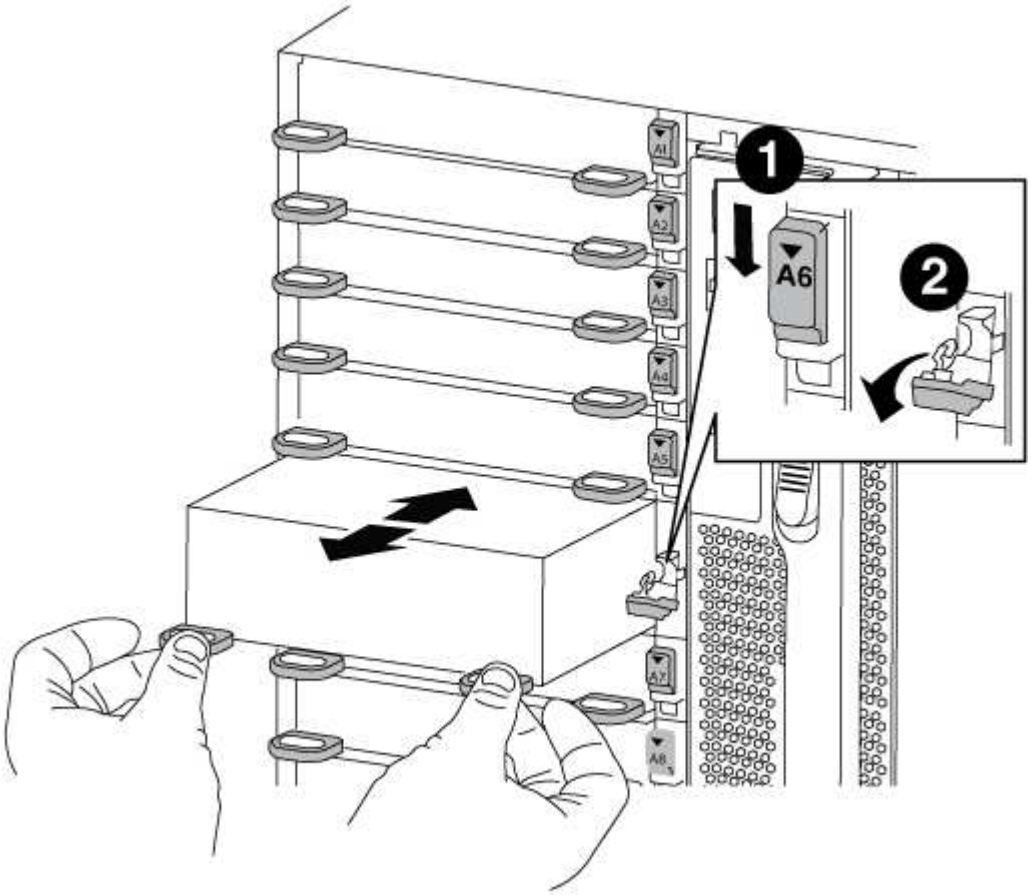

1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

4. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

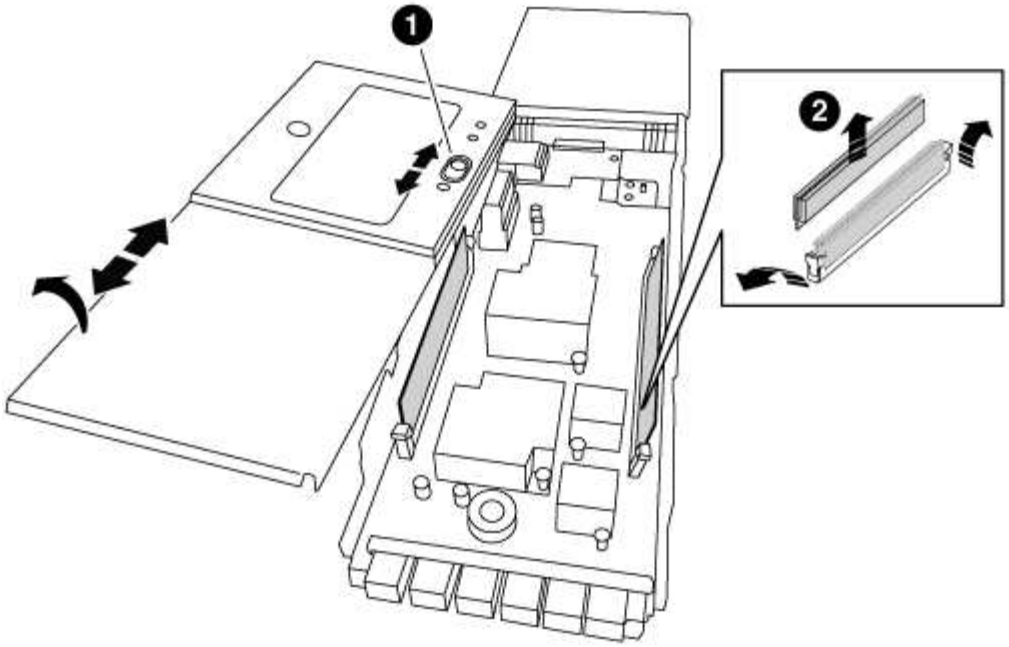

1	Cover locking button
2	DIMM and DIMM ejector tabs

5. Remove the DIMMs, one at a time, from the old NVRAM module and install them in the replacement NVRAM module.
6. Close the cover on the module.
7. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

Step 3: Replace a NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, you must remove the NVRAM module, open the module, and then replace the target DIMM.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the target NVRAM module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The NVRAM module disengages from the chassis and moves out a few inches.

   c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module

face.

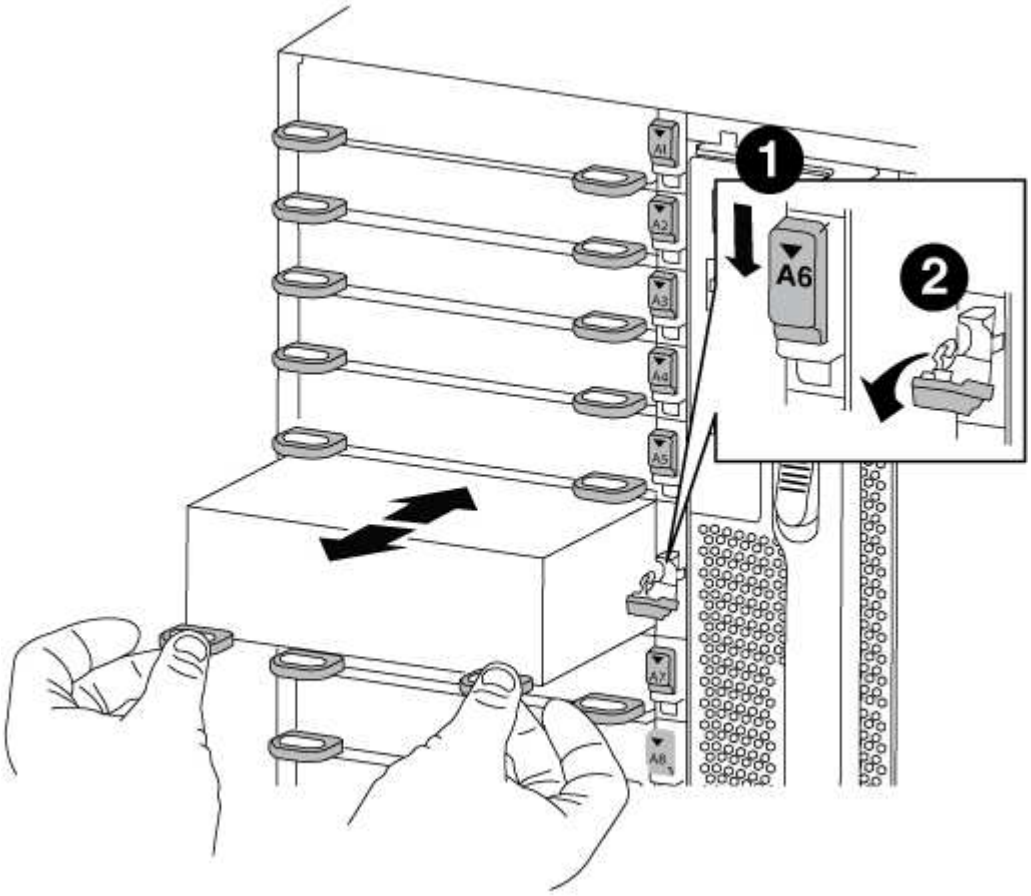

1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.

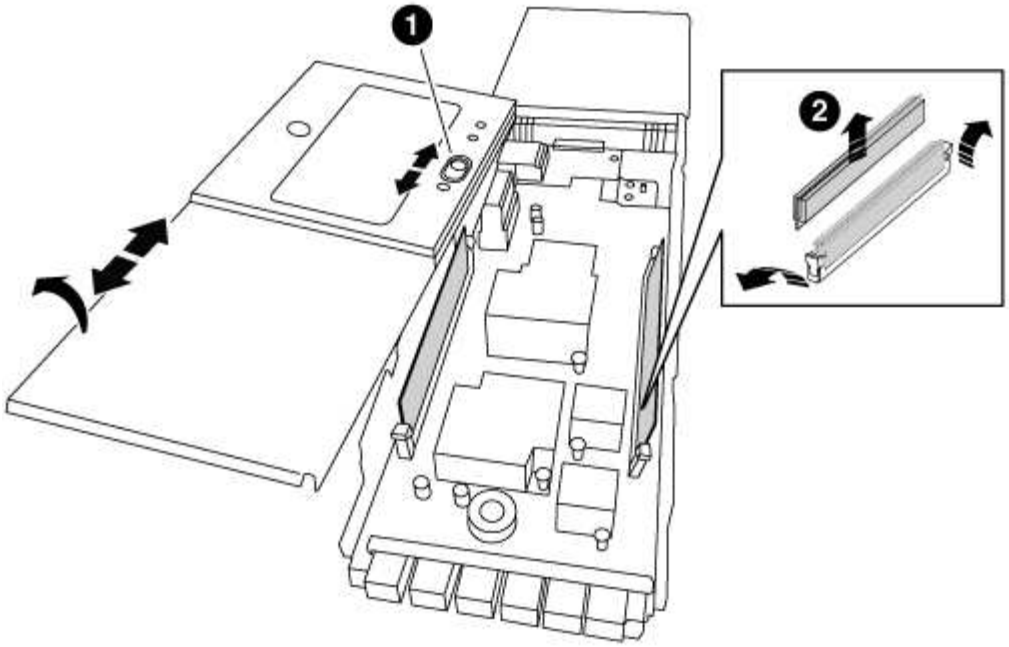

1	Cover locking button
2	DIMM and DIMM ejector tabs

4. Locate the DIMM to be replaced inside the NVRAM module, and then remove it by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
5. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
6. Close the cover on the module.
7. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

Step 4: Reboot the controller after FRU replacement

After you replace the FRU, you must reboot the controller module.

Step

1. To boot ONTAP from the LOADER prompt, enter `bye`.

Step 5: Reassign disks

Depending on whether you have an HA pair or two-node MetroCluster configuration, you must either verify the reassignment of disks to the new controller module or manually reassign the disks.

Select one of the following options for instructions on how to reassign disks to the new controller.

Option 1: Verify ID (HA pair)

Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* node and then verify that the change was implemented.

> Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the replacement node is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the replacement node, boot the node, entering `y` if you are prompted to override the system ID due to a system ID mismatch.

   ```
   boot_ontap bye
   ```

   The node will reboot, if autoboot is set.

3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* node console and then, from the healthy node, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired node, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

   ```
   node1> `storage failover show`
                                       Takeover
   Node              Partner           Possible     State Description
   ------------      ------------      --------
   -------------------------------------
   node1             node2             false        System ID changed
   on partner (Old:
                                                     151759755, New:
   151759706), In takeover
   node2             node1             -            Waiting for
   giveback (HA mailboxes)
   ```

4. From the healthy node, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`

   c. Wait for the `savecore`command to complete before issuing the giveback.

You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

d. Return to the admin privilege level: `set -privilege admin`

5. Give back the node:

a. From the healthy node, give back the replaced node's storage: `storage failover giveback -ofnode replacement_node_name`

The *replacement* node takes back its storage and completes booting.

If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

If the giveback is vetoed, you can consider overriding the vetoes.

Find the High-Availability Configuration Guide for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the `System ID changed on partner` message.

6. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* node should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home
ID Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

7. If the system is in a MetroCluster configuration, monitor the status of the node: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each node will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

8. If the node is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a node on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The *replacement* node is the current owner of the disks on the disaster site.

     Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

9. If your system is in a MetroCluster configuration, verify that each node is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

10. Verify that the expected volumes are present for each node: `vol show -node node-name`
11. If you disabled automatic takeover on reboot, enable it from the healthy node: `storage failover modify -node replacement-node-name -onreboot true`

Option 2: Reassign ID (MetroCluster config)

Reassign the system ID in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `metrocluster node show -fields node-systemid,dr-partner-systemid`

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

```
dr-group-id cluster         node                 node-systemid dr-
partner-systemid
 ----------- --------------------- --------------------
------------ -------------------
 1           Cluster_A             Node_A_1             536872914
118073209
 1           Cluster_B             Node_B_1             118073209
536872914
 2 entries were displayed.
```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

```
Local System ID: 118065481
    ...
    ...
```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                  POOL   SERIAL NUMBER  HOME
-------    -------------          -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Verify that the coredumps are saved: `system node run -node` *`local-node-name`* `partner savecore`

      If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node` *`local-node-name`* `partner savecore -s command`.</info>.

   c. Return to the admin privilege level: `set -privilege admin`
7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
8. Boot the *replacement* node: `boot_ontap`
9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`
10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

       You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

    b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`
    c. Return to the admin privilege level: `set -privilege admin`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - FAS9000

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

All other components in the system must be functioning properly; if not, you must contact technical support.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- The number of power supplies in the system depends on the model.
- Power supplies are auto-ranging.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Press and hold the orange button on the power supply handle, and then pull the power supply out of the chassis.

> When removing a power supply, always use two hands to support its weight.

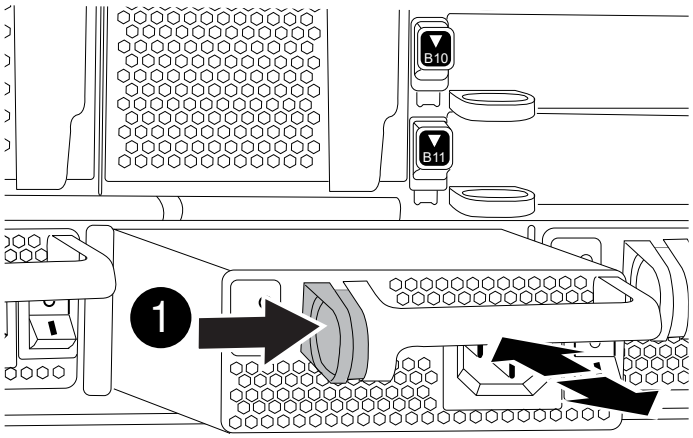

1	Locking button

5. Make sure that the on/off switch of the new power supply is in the Off position.

6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

   The power supplies are keyed and can only be installed one way.

   > Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply and the power source.

   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.

8. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

   The green power LED lights when the PSU is fully inserted into the chassis and the amber attention LED flashes initially, but turns off after a few moments.

9. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - AFF 9000

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node local -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the orange button on the cam handle downward until it unlocks.

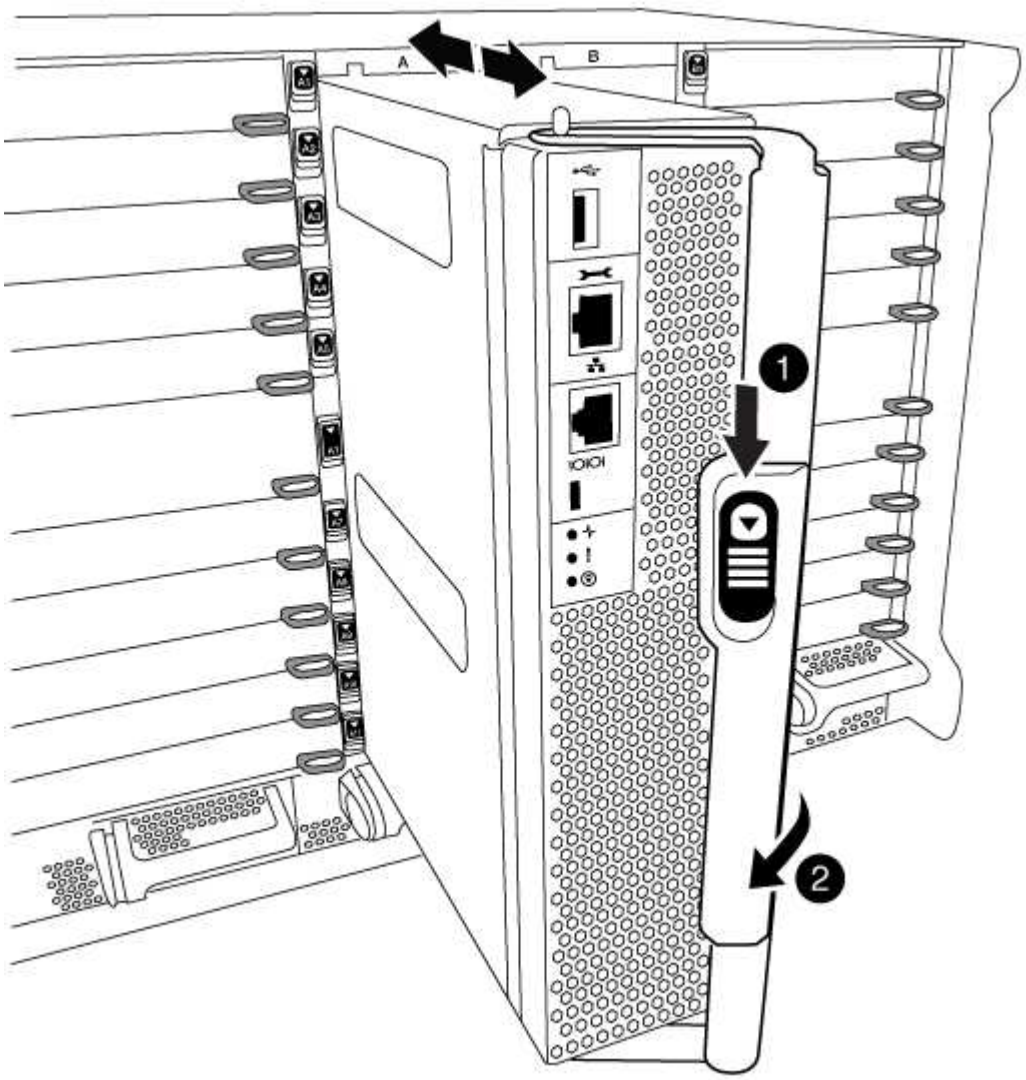

1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.

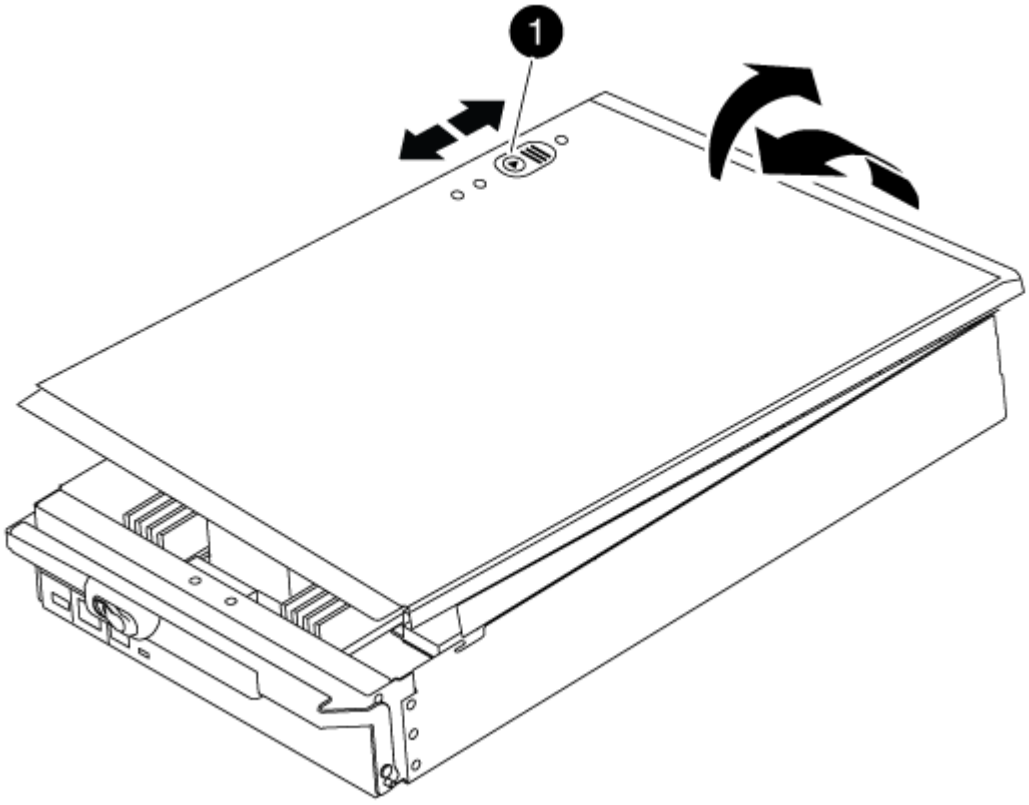

1	Controller module cover locking button

Step 3: Replace the RTC battery

To replace the RTC battery, you must locate the failed battery in the controller module, remove it from the holder, and then install the replacement battery in the holder.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

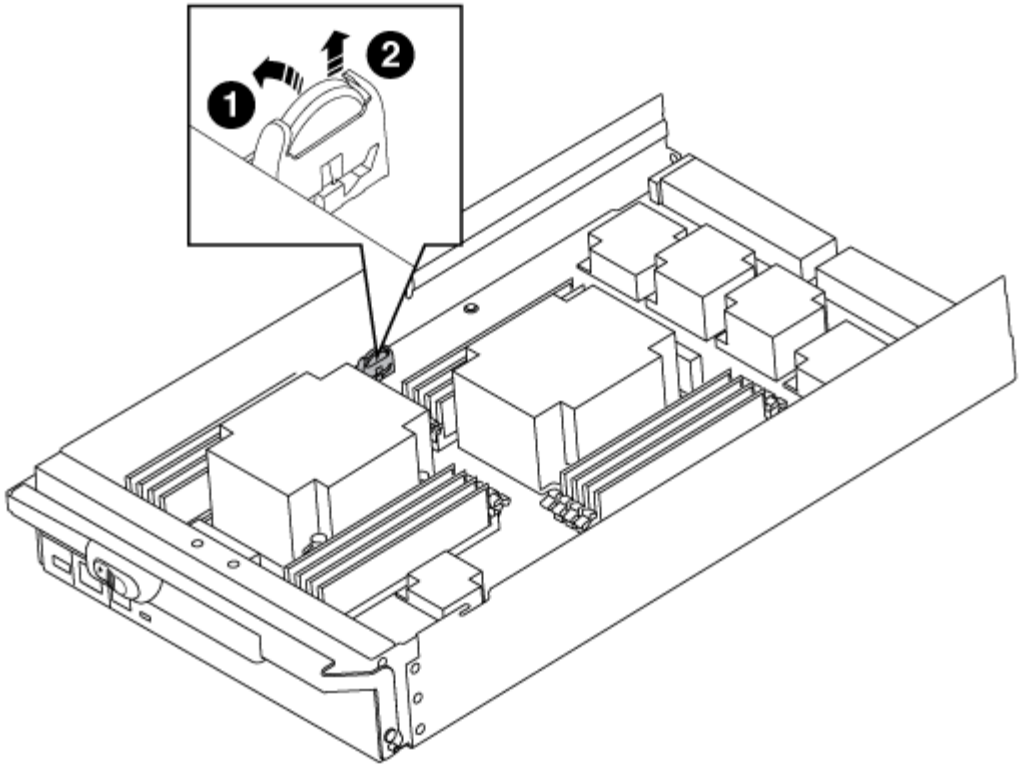

1	RTC battery
2	RTC battery housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

> Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

4. Remove the replacement battery from the antistatic shipping bag.
5. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
8. Reinstall the controller module cover.

Step 4: Reinstall the controller module and set time/date

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

Steps

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy node with the `show date` command.
   b. At the LOADER prompt on the target node, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target node.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the node reboot.
8. Return the node to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

X91148A module

Overview of adding an X91148A module - AFF A9000

You can add an I/O module to your system by either replacing a NIC or storage adapter with a new one in a fully-populated system, or by adding a new NIC or storage adapter into an empty chassis slot in your system.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- To non-disruptively add an I/O module, you must takeover the target controller, remove the slot blanking cover in the target slot or remove an existing I/O module, add the new or replacement I/O module, and then giveback the target controller.
- Make sure that all other components are functioning properly.

Add an X91148A module in a system with open slots - FAS9000

You can add an X91148A module into an empty module slot in your system as either a 100GbE NIC or a storage module for the NS224 storage shelves.

- Your system must be running ONTAP 9.8 and later.
- To non-disruptively add the X91148A module, you must takeover the target controller, remove the slot blanking cover in the target slot, add the module, and then giveback the target controller.
- There must be one or more open slots available on your system.
- If multiple slots are available, install the module according to the slot priority matrix for the X91148A module in the NetApp Hardware Universe.
- If you are adding the X91148A module as a storage module, you must install the module slots 3 and/or 7.
- If you are adding the X91148A module as a 100GbE NIC, you can use any open slot. However, by default, slots 3 and 7 are set as storage slots. If you wish to use those slots as network slots and will not add NS224 shelves, you must modify the slots for networking use with the `storage port modify -node` *`node name`* `-port` *`port name`* `-mode network` command. See the NetApp Hardware Universe for other slots that can be used by the X91148A module for networking.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Option 1: Add an X91148A module as a NIC module

To add an X91148A module as a NIC module in a system with open slots, you must follow the specific sequence of steps.

Steps

1. Shutdown controller A:
   a. Disable automatic giveback: `storage failover modify -node local -auto-giveback false`
   b. Take over the target node: `storage failover takeover -ofnode target_node_name`

      The console connection shows that the node drops to the LOADER prompt when the takeover is complete.
2. If you are not already grounded, properly ground yourself.
3. Remove the target slot blanking cover:
   a. Depress the lettered and numbered cam button.
   b. Rotate the cam latch down until it is in a horizontal position.
   c. Remove the blanking cover.
4. Install the X91148A module:
   a. Align the X91148A module with the edges of the slot.
   b. Slide the X91148A module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
5. Cable the module to the data switches.
6. Reboot controller A from the LOADER prompt: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.
7. Giveback the node from the partner node: `storage failover giveback -ofnode target_node_name`
8. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto-giveback true`
9. Repeat these steps for controller B.

Option 2: Add an X91148A module as a storage module

To add an X91148A module as a storage module in a system with open slots, you must follow the specific sequence of steps.

- This procedure presumes slots 3 and/or 7 are open.

Steps

1. Shut down controller A:
   a. Disable automatic giveback: `storage failover modify -node local -auto-giveback false`

b. Take over the target node: `storage failover takeover -ofnode target_node_name`

   The console connection shows that the node drops to the LOADER prompt when the takeover is complete.

2. If you are not already grounded, properly ground yourself.
3. Remove the target slot blanking cover:
   a. Depress the lettered and numbered cam button.
   b. Rotate the cam latch down until it is in a horizontal position.
   c. Remove the blanking cover.
4. Install the X91148A module into slot 3:
   a. Align the X91148A module with the edges of the slot.
   b. Slide the X91148A module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
   d. If you are installing a second X91148A module for storage, repeat this step for the module in slot 7.
5. Reboot controller A:
   - If the replacement module is not the same model as the old module, reboot the BMC :
     a. From the LOADER prompt, change to advanced privilege mode: `set -privilege advanced`
     b. Reboot the BMC: `sp reboot`
   - If the replacement module is the same as the old module, boot from the LOADER prompt: `bye`

     This reinitializes the PCIe cards and other components and reboots the node.
6. Giveback the node from the partner node: `storage failover giveback -ofnode target_node_name`
7. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto-giveback true`
8. Repeat these steps for controller B.
9. Install and cable your NS224 shelves, as described in Hot-add workflow.

Add an X91148A storage module in a system with no open slots - FAS9000

You must remove one more or more existing NIC or storage modules in your system in order to install one or more X91148A storage modules into your fully-populated system.

- Your system must be running ONTAP 9.8 and later.
- To non-disruptively add the X91148A module, you must takeover the target controller, add the module, and then giveback the target controller.
- If you are adding the X91148A module as a storage adapter, you must install the module in slots 3 and/or 7.

- If you are adding the X91148A module as a 100GbE NIC, you can use any open slot. However, by default, slots 3 and 7 are set as storage slots. If you wish to use those slots as network slots and will not add NS224 shelves, you must modify the slots for networking use with the `storage port modify -node` *`node name`* `-port` *`port name`* `-mode network` command for each port. See the NetApp Hardware Universe for other slots that can be used by the X91148A module for networking.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Option 1: Add an X91148A module as a NIC module

You must remove one or more existing NIC or storage modules in your system in order to install one or more X91148A NIC modules into your fully-populated system.

Steps

1. If you are adding an X91148A module into a slot that contains a NIC module with the same number of ports as the X91148A module, the LIFs will automatically migrate when its controller module is shut down. If the NIC module being replaced has more ports than the X91148A module, you must permanently reassign the affected LIFs to a different home port. See Migrating a LIF for information about using System Manager to permanently move the LIFs
2. Shut down controller A:
   a. Disable automatic giveback: `storage failover modify -node local -auto-giveback false`
   b. Take over the target node: `storage failover takeover -ofnode target_node_name`

      The console connection shows that the node drops to the LOADER prompt when the takeover is complete.
3. If you are not already grounded, properly ground yourself.
4. Unplug any cabling on the target I/O module.
5. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.
   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.
   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

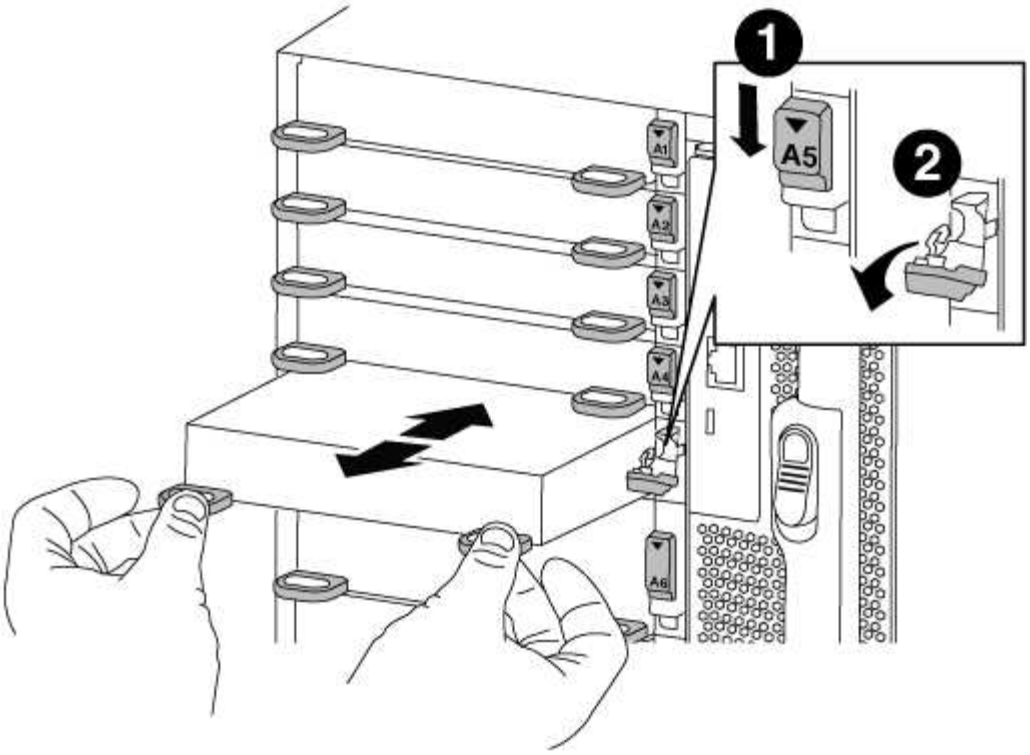

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

6. Install the X91148A module into the target slot:
   a. Align the X91148A module with the edges of the slot.
   b. Slide the X91148A module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
7. Repeat the remove and install steps to replace additional modules for controller A.
8. Cable the module or modules to the data switches.
9. Reboot controller A from the LOADER prompt: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.

10. Giveback the node from the partner node: `storage failover giveback -ofnode target_node_name`
11. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto -giveback true`
12. If you added the X91148A module as a NIC module in slots 3 or 7, for networking, use the `storage port modify -node` *node name* `-port` *port name* `-mode network` command for each port.

13. Repeat these steps for controller B.

Option 2: Adding an X91148A module as a storage module

You must remove one or more existing NIC or storage modules in your system in order to install one or more X91148A storage modules into your fully-populated system.

- This procedure presumes you re installing the X91148A module into slots 3 and/or 7.

Steps

1. If you are adding an X91148A module as a storage module in slots 3 and/or 7 into a slot that has an existing NIC module in it, use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.
2. Shut down controller A:
   a. Disable automatic giveback: `storage failover modify -node local -auto-giveback false`
   b. Take over the target node: `storage failover takeover -ofnode target_node_name`

      The console connection shows that the node drops to the LOADER prompt when the takeover is complete.
3. If you are not already grounded, properly ground yourself.
4. Unplug any cabling on the target I/O module.
5. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.
   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.
   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

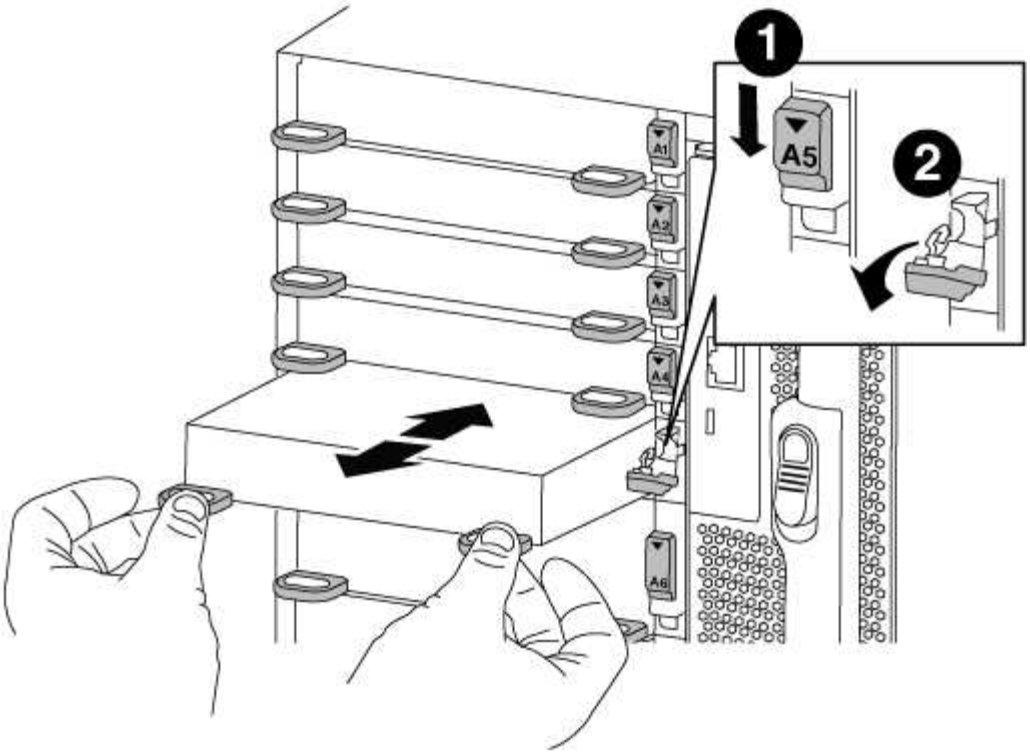

1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

6. Install the X91148A module into slot 3:
   a. Align the X91148A module with the edges of the slot.
   b. Slide the X91148A module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
   d. If you are installing a second X91148A module for storage, repeat the remove and install steps for the module in slot 7.
7. Reboot controller A from the LOADER prompt: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.

8. Giveback the node from the partner node: `storage failover giveback -ofnode` *`target_node_name`*
9. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto -giveback true`
10. Repeat these steps for controller B.
11. Install and cable your NS224 shelves, as described in Hot-add workflow.

Other models

Documentation for AFF and FAS systems that have reached the end of hardware support are available for customer use in the Archive. Documentation for older AFF and FAS models that are no longer available for purchase, but are still supported are available in the A-Z Documentation Library.

Drive shelves for ONTAP hardware systems

NS224 shelves

Hot-add shelf

Hot-add workflow - NS224 shelves

Follow these workflow steps to hot-add your NS224 shelf.

Before you begin

- This procedure applies to direct-attached storage only. To view instructions for switch-attached storage, view our switch-attached cabling guide.
- To hot-add an NS224 shelf, your HA pair must meet certain requirements. Review the hot-add requirements and best practices.

Prepare to hot-add your shelf

Depending on your platform model, you might need to install additional RoCE-capable PCIe cards or I/O modules, configure the non-dedicated RoCE-capable Ethernet ports for storage use, recable an existing shelf across two sets of ports in different slots for resiliency against slot failure, and disable automatic drive assignment if you are manually assigning drive ownership.

Install your shelf

To install the shelf, install the rail kit for the shelf, and then install and secure your shelf in the telco rack or cabinet. Next, connect the power cords to power on the shelf and then assign a unique shelf ID to ensure the shelf is distinct within the HA pair.

Cable your shelf

Cable the shelf you are hot-adding so that it has two connections to each controller in the HA pair.

Complete your hot-add

If you disabled automatic drive assignment as part of the preparation for the hot-add, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed.

Requirements and best practices to hot-add NS224 shelves

Before hot-adding a shelf, be sure to review the requirements and best practices.

Requirements

To hot-add an NS224 shelf, your HA pair must meet certain requirements.

- **Supported ONTAP version**: Your platform model and version of ONTAP must support the NS224 shelf and drives you are hot-adding. See NetApp Hardware Universe

- **Number of shelves**: Your HA pair must have less than the maximum number of shelves supported, by at least the number of shelves you plan to hot-add.

  You cannot have exceeded the maximum number of shelves supported by your HA pair after hot-adding shelves. See NetApp Hardware Universe.

- **Cabling**:
  - Make sure you have the correct number and type of cables to connect the shelf. See NetApp Hardware Universe.
  - If you are hot-adding a shelf to an HA pair that already has an NS224 shelf, your HA pair cannot have any storage cabling error messages, and it must be cabled as multipath HA.

    You can run Active IQ Config Advisor to view any storage cabling error messages and the corrective actions you should take.

Best practices

Familiarize yourself with the following best practices before hot-adding an NS224 shelf.

- **Disk Qualification Package:** The best practice is to have the current version of the Disk Qualification Package installed before hot-adding a shelf.

  Having the current version of the DQP installed allows your system to recognize and use newly qualified drives. This avoids system event messages about having noncurrent drive information and prevention of drive partitioning because drives are not recognized. The DQP also notifies you of noncurrent drive firmware.

- **Active IQ Config Advisor:** The best practice is to run Active IQ Config Advisor before and after hot-adding a shelf.

  Running Active IQ Config Advisor before hot-adding a shelf provides a snapshot of the existing shelf Ethernet (ENET) connectivity, verifies NVMe shelf module (NSM) firmware versions, and allows you to verify a shelf ID already in use in the HA pair.

  Running Active IQ Config Advisor after hot-adding a shelf allows you to verify shelves are cabled correctly and that shelf IDs are unique within the HA pair.

- **NSM firmware:** The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your storage system before adding a new shelf.

  Do not revert firmware to a version that does not support your shelf and its components.

Prepare for a hot-add - NS224 shelves

Complete the preparation tasks applicable to your HA pair before hot-adding an NS224 shelf.

When you are done with the applicable preparation tasks, go to Install a shelf for a hot-add.

Install RoCE-capable PCIe cards or I/O modules

If your platform model supports the use of RoCE-capable PCIe cards or I/O modules, your HA pair must have

enough available RoCE-capable Ethernet ports to support the number of shelves you are hot-adding.

Steps

1. For each shelf you are hot-adding, verify that there are two RoCE-capable ports on each controller.

   These ports can be on board the controllers, on RoCE-capable PCIe cards, a combination of both, or on RoCE-capable I/O modules, as supported by your platform model.

2. If your HA pair does not have enough available RoCE-capable ports, install the additional PCIe cards or I/O modules into the correct controller slots, as supported by your platform model.
   a. Identify the correct controller slots for your platform model. See NetApp Hardware Universe.
   b. See your platform model's documentation for PCIe card or I/O module installation instructions.

Configure RoCE-capable ports for storage use

If your HA pair has non-dedicated RoCE-capable Ethernet ports that you are using to hot-add an NS224 shelf, you must configure the ports for storage use (not networking use).

Before you begin

- Make sure you installed any additional RoCE-capable PCIe cards or I/O modules in each controller.

About this task

- For some platform models, when a RoCE-capable PCIe card or I/O module is installed in a supported slot on a controller, the ports automatically default to storage use (instead of networking); however, it is recommended that you complete the following procedure to verify the RoCE-capable ports are configured for storage use.
- If you determine that the non-dedicated RoCE-capable ports in your HA pair are not configured for storage use, it is a nondisruptive procedure to configure them. You do not need to reboot the controllers, unless one or both controllers are in maintenance mode. This procedure assumes that neither controller is in maintenance mode.
- If in the future you need to change ports from storage use to networking use, enter the command, `storage port modify -node node_name -port port_name -mode network`.

Steps

1. Log in to the cluster using SSH or the serial console port.
2. Enter the following command to verify if the non-dedicated ports in the HA pair are configured for storage use:

   ```
   storage port show
   ```

   - If your HA pair is running ONTAP 9.8 or later, the non-dedicated ports display `storage` in the `Mode` column.
   - If your HA pair is running ONTAP 9.7, the non-dedicated ports, which display `false` in the `Is Dedicated?` column, also display `enabled` in the `State` column.

When non-dedicated ports are not configured for storage use, the command output displays the following:

- If your HA pair is running ONTAP 9.8 or later, the non-dedicated ports display `network` in the `Mode` column.
- If your HA pair is running ONTAP 9.7, the non-dedicated ports, which display `false` in the `Is Dedicated?` column, also display `disabled` in the `State` column.

3. If the non-dedicated ports are configured for storage use, you are done with this procedure.

   Otherwise, you need to configure the ports by completing steps 4 through 7.

4. Configure the non-dedicated ports for storage use, on one of the controllers:

   You must repeat the applicable command for each port you are configuring.

If your HA pair is running…	Use this command…
ONTAP 9.8 or later	`storage port modify -node` *`node_name`* `-port` *`port_name`* `-mode storage`
ONTAP 9.7	`storage port enable -node` *`node_name`* `-port` *`port_name`*

5. Repeat step 4 for the second controller.
6. Verify that the non-dedicated ports on both controllers are configured for storage use: `storage port show`
   - If your HA pair is running ONTAP 9.8 or later, the non-dedicated ports display `storage` in the `Mode` column.
   - If your HA pair is running ONTAP 9.7, the non-dedicated ports, which display `false` in the `Is Dedicated?` column, also display `enabled` in the `State` column.

Recable existing shelves

Before hot-adding additional shelves, depending on your platform model, you might need to recable an existing shelf (after you have installed the additional RoCE-capable PCIe cards or I/O modules) across two sets of ports in different slots for resiliency against slot failure.

Before you begin

- Make sure you installed any additional RoCE-capable PCIe cards or I/O modules in each controller.
- Make sure that the non-dedicated ports on the RoCE-capable PCIe cards or I/O modules you installed are configured for storage use.

About this task

- Recabling port connections is a nondisruptive procedure when your shelf has multipath-HA connectivity.
- You move one cable at a time to always maintain connectivity to the shelf during this procedure.

Moving a cable does not require any wait time between unplugging the cable from one port and plugging it into another port.

- If needed, refer to the shelf cabling illustrations for your platform model in Overview of cabling for a hot-add.

Steps

1. Recable the existing shelf's connections across two sets of RoCE-capable ports in different slots, as applicable to your platform model.

- For AFF systems:

AFF A1K

Do one of the following if you are hot-adding a second shelf or a fourth shelf.

> If you have an AFF A1K HA pair and you are hot-adding a third shelf and installing a third or fourth RoCE-capable I/O module in each controller, the third shelf is cabled to only the third or third and fourth I/O modules. You do not need to recable any existing shelves.

- If you are hot-adding a second shelf, recable the first shelf across the RoCE-capable I/O modules in slot 11 and slot 10 on each controller.

  The substeps assume the existing shelf is cabled to a RoCE-capable I/O module in slot 11 on each controller.

  a. On controller A, move the cable from slot 11 port b (e11b) to slot 10 port b (e10b).

  b. Repeat the same cable move on controller B.

- If you are hot-adding a fourth shelf, recable the third shelf across the RoCE-capable I/O modules in slot 9 and slot 8 on each controller.

  The substeps assume the third shelf is cabled to a RoCE-capable I/O module in slot 9 on each controller.

  a. On controller A, move the cable from slot 9 port b (e9b) to slot 8 port b (e8b).

  b. Repeat the same cable move on controller B.

AFF A70, AFF A90, or AFF C80

If you are hot-adding a second shelf, recable the first shelf across the RoCE-capable I/O modules in slot 11 and slot 8 on each controller.

The substeps assume the existing shelf is cabled to a RoCE-capable I/O module in slot 11 on each controller.

a. On controller A, move the cable from slot 11 port b (e11b) to slot 8 port b (e8b).

b. Repeat the same cable move on controller B.

AFF A800 or AFF C800

If you are hot-adding a second shelf, recable the first shelf across the two sets of RoCE-capable ports in slot 5 and slot 3 on each controller.

The substeps assume the existing shelf is cabled to RoCE-capable PCIe cards in slot 5 on each controller.

a. On controller A, move the cable from slot 5 port b (e5b) to slot 3 port b (e3b).

b. Repeat the same cable move on controller B.

AFF A700

If you are hot-adding a second shelf, recable the first shelf across the two sets of RoCE-capable ports

in slot 3 and slot 7 on each controller.

The substeps assume the existing shelf is cabled to RoCE-capable I/O modules in slot 3 on each controller.

a. On controller A, move the cable from slot 3 port b (e3b) to slot 7 port b (e7b).
b. Repeat the same cable move on controller B.

AFF A400 or AFF C400

If you are hot-adding a second shelf, depending on your platform model, do one of the following:

- On the AFF A400:

  Recable the first shelf across the two sets of RoCE-capable ports, onboard e0c/e0d and in slot 5, on each controller.

  The substeps assume the existing shelf is cabled to RoCE-capable onboard ports e0c/e0d on each controller.

  a. On controller A, move the cable from port e0d to slot 5 Port b (e5b).
  b. Repeat the same cable move on controller B.

- On the AFF C400:

  Recable the first shelf across the two sets of RoCE-capable ports in slot 4 and slot 5, on each controller.

  The substeps assume the existing shelf is cabled to RoCE-capable ports in slot 4 on each controller.

  a. On controller A, move the cable from slot 4 port a (e4a) to slot 5 port b (e5b).
  b. Repeat the same cable move on controller B.

AFF A900

Do one of the following if you are hot-adding a second shelf or a fourth shelf.

- If you are hot-adding a second shelf, recable the first shelf across the RoCE-capable I/O modules in slot 2 and slot 10 on each controller.

  The substeps assume the existing shelf is cabled to a RoCE-capable I/O module in slot 2 on each controller.

  a. On controller A, move the cable from slot 2 port b (e2b) to slot 10 port b (e10b).
  b. Repeat the same cable move on controller B.

- If you are hot-adding a fourth shelf, recable the third shelf across the RoCE-capable I/O modules in slot 1 and slot 11 on each controller.

  The substeps assume the third shelf is cabled to a RoCE-capable I/O module in slot 1 on each controller.

  a. On controller A, move the cable from slot 1 port b (e1b) to slot 11 port b (e11b).

b. Repeat the same cable move on controller B.

AFF A30, AFF C30, AFF A50, or AFF C60

If you are hot-adding a second shelf, recable the first shelf across the RoCE-capable I/O modules in slot 3 and slot 1 on each controller.

The substeps assume the existing shelf is cabled to a RoCE-capable I/O module in slot 3 on each controller.

a. On controller A, move the cable from slot 3 port b (e3b) to slot 1 port b (e1b).

b. Repeat the same cable move on controller B.

- For ASA systems:

ASA A1K

Do one of the following if you are hot-adding a second shelf or a fourth shelf.

> If you have an ASA A1K HA pair and you are hot-adding a third shelf and installing a third or fourth RoCE-capable I/O module in each controller, the third shelf is cabled to only the third or third and fourth I/O modules. You do not need to recable any existing shelves.

- If you are hot-adding a second shelf, recable the first shelf across the RoCE-capable I/O modules in slot 11 and slot 10 on each controller.

  The substeps assume the existing shelf is cabled to a RoCE-capable I/O module in slot 11 on each controller.

  a. On controller A, move the cable from slot 11 port b (e11b) to slot 10 port b (e10b).

  b. Repeat the same cable move on controller B.

- If you are hot-adding a fourth shelf, recable the third shelf across the RoCE-capable I/O modules in slot 9 and slot 8 on each controller.

  The substeps assume the third shelf is cabled to a RoCE-capable I/O module in slot 9 on each controller.

  a. On controller A, move the cable from slot 9 port b (e9b) to slot 8 port b (e8b).

  b. Repeat the same cable move on controller B.

ASA A70 or ASA A90

If you are hot-adding a second shelf, recable the first shelf across the RoCE-capable I/O modules in slot 11 and slot 8 on each controller.

The substeps assume the existing shelf is cabled to a RoCE-capable I/O module in slot 11 on each controller.

a. On controller A, move the cable from slot 11 port b (e11b) to slot 8 port b (e8b).

b. Repeat the same cable move on controller B.

ASA A800 or ASA C800

If you are hot-adding a second shelf, recable the first shelf across the two sets of RoCE-capable ports in slot 5 and slot 3 on each controller.

The substeps assume the existing shelf is cabled to RoCE-capable PCIe cards in slot 5 on each controller.

a. On controller A, move the cable from slot 5 port b (e5b) to slot 3 port b (e3b).

b. Repeat the same cable move on controller B.

ASA A400 or ASA C400

If you are hot-adding a second shelf, depending on your platform model, do one of the following:

- On the ASA A400:

  Recable the first shelf across the two sets of RoCE-capable ports, onboard e0c/e0d and in slot 5, on each controller.

  The substeps assume the existing shelf is cabled to RoCE-capable onboard ports e0c/e0d on each controller.

  a. On controller A, move the cable from port e0d to slot 5 Port b (e5b).
  b. Repeat the same cable move on controller B.

- On the ASA C400:

  Recable the first shelf across the two sets of RoCE-capable ports in slot 4 and slot 5, on each controller.

  The substeps assume the existing shelf is cabled to RoCE-capable ports in slot 4 on each controller.

  a. On controller A, move the cable from slot 4 port a (e4a) to slot 5 port b (e5b).
  b. Repeat the same cable move on controller B.

ASA A900

Do one of the following if you are hot-adding a second shelf or a fourth shelf.

- If you are hot-adding a second shelf, recable the first shelf across the RoCE-capable I/O modules in slot 2 and slot 10 on each controller.

  The substeps assume the existing shelf is cabled to a RoCE-capable I/O module in slot 2 on each controller.

  a. On controller A, move the cable from slot 2 port b (e2b) to slot 10 port b (e10b).
  b. Repeat the same cable move on controller B.

- If you are hot-adding a fourth shelf, recable the third shelf across the RoCE-capable I/O modules in slot 1 and slot 11 on each controller.

  The substeps assume the third shelf is cabled to a RoCE-capable I/O module in slot 1 on each controller.

  a. On controller A, move the cable from slot 1 port b (e1b) to slot 11 port b (e11b).
  b. Repeat the same cable move on controller B.

ASA A30 or ASA A50

If you are hot-adding a second shelf, recable the first shelf across the RoCE-capable I/O modules in slot 3 and slot 1 on each controller.

The substeps assume the existing shelf is cabled to a RoCE-capable I/O module in slot 3 on each controller.

a. On controller A, move the cable from slot 3 port b (e3b) to

slot 1 port b (e1b).

b. Repeat the same cable move on controller B.

2. Verify that the recabled shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

Disable automatic drive assignment

If you are manually assigning drive ownership for the NS224 shelf you are hot-adding, then you need to disable automatic drive assignment if it is enabled.

If you are unsure whether you should manually assign drive ownership, or want to understand the automatic assignment of drive ownership policies for your storage system, go to About automatic assignment of disk ownership.

Steps

1. Verify whether automatic drive assignment is enabled: `storage disk option show`

   You can enter the command on either node.

   If automatic drive assignment is enabled, the output shows `on` in the `Auto Assign` column (for each node).

2. If automatic drive assignment is enabled, disable it: `storage disk option modify -node` *`node_name`* `-autoassign off`

   You must disable automatic drive assignment on both nodes.

Install a shelf for a hot-add - NS224 shelves

You need to install an NS224 shelf in a cabinet or telco rack, connect the power cords (which automatically powers on the shelf), and set the shelf ID.

Before you begin

- Make sure your have a paper clip with one side straightened or a narrow-tipped ballpoint pen.

  To change the shelf ID, you use the paper clip or ballpoint pen to access the shelf ID button behind the Operator Display Panel (ODP) for the change shelf ID step.

- Understand that a fully loaded NS224 shelf can weigh up to 66.78 lbs (30.29 kg) with NSM100 modules or an average of 56.8 lbs (25.8 kg) with NSM100B modules and requires two people to lift or use of a hydraulic lift. Avoid removing shelf components (from the front or rear of the shelf) to reduce the shelf weight, because shelf weight will become unbalanced.

Steps

1. Install the rail kit for the shelf, as needed, using the instructions included with the kit.

   Always use the appropriate rail kit for your shelf to install the shelf in a rack or cabinet.

2. Install the shelf:
   a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

      If you are installing multiple shelves, place the first shelf directly above the controllers. Place the second shelf directly under the controllers. Repeat this pattern for any additional shelves.

   b. Secure the shelf to the cabinet or telco rack using the mounting screws included in the kit.
3. Connect the power:
   a. Connect the power cords to the shelf and secure them in place.

      If they are AC power supplies, secure them in place with the power cord retainer.

      If they are DC power supplies, secure them in place with the two thumb screws.

   b. Connect the power cords to different power sources for resiliency.

      A shelf powers up when connected to a power source; it does not have power switches. When functioning correctly, a power supply's bicolored LED illuminates green.
4. Set the shelf ID to a number that is unique within the HA pair:

   For more detailed instructions, see Change a shelf ID - NS224 shelves.

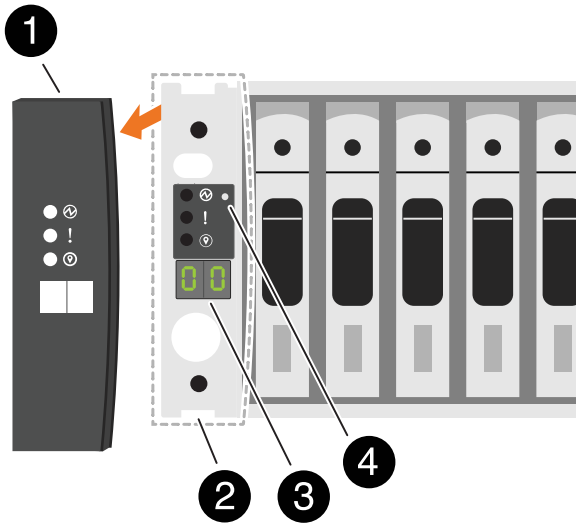

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID number
4	Shelf ID button

   a. Remove the left end cap and locate the small hole to the right of the LEDs.

b. Insert the end of a paper clip or similar tool into the small hole to reach the shelf ID button.

c. Press and hold the button (for up to 15 seconds) until the first number on the digital display blinks, and then release the button.

   If the ID takes longer than 15 seconds to blink, press and hold the button again, making sure to press it in all the way.

d. Press and release the button to advance the number until you reach the desired number from 0 to 9.

e. Repeat substeps 4c and 4d to set the second number of the shelf ID.

   It can take up to three seconds (instead of 15 seconds) for the number to blink.

f. Press and hold the button until the second number stops blinking.

   After about five seconds, both numbers start blinking and the amber LED on the ODP illuminates.

g. Power-cycle the shelf to make the shelf ID take effect.

   You must unplug both power cords from the shelf, wait 10 seconds, and then plug them back in.

   When power is restored to the power supplies, their bicolored LEDs illuminate green.

What's next?

Cable your hot-add shelf. Go to Overview of cabling for a hot-add.

Cable shelf for hot-add

Overview of cabling for a hot-add - NS224 shelves

You cable each NS224 shelf you are hot-adding so that each shelf has two connections to each controller in the HA pair.

This cabling section describes how to cable the NS224 shelf to the following storage systems:

- Cable to AFF systems
- Cable to ASA systems
- Cable to EOA systems

About this task

- This procedure applies to direct-attached storage only. To view instructions for switch-attached storage, view our switch-attached cabling guide.
- Familiarize yourself with proper cable connector orientation, and the location and labeling of ports on the NS224 NSM100 shelf modules.
  - Cables are inserted with the connector pull-tab facing up.

    When a cable is inserted correctly, it clicks into place.

    After you connect both ends of the cable, the shelf and controller port LNK (green) LEDs illuminate. If a port LNK LED does not illuminate, reseat the cable.

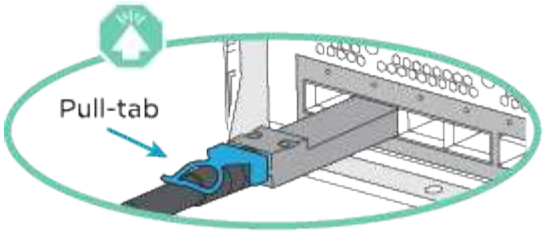

◦ You can use the following illustration to help you physically identify the shelf NSM100 ports, e0a and e0b.

NSM100 modules

- An NS224 shelf contains two NSM100 modules. The top module goes in slot A (NSM A) and the bottom module goes in slot B (NSM B).
- Each NSM100 module includes 2 x 100GbE QSFP28 ports: e0a and e0b.

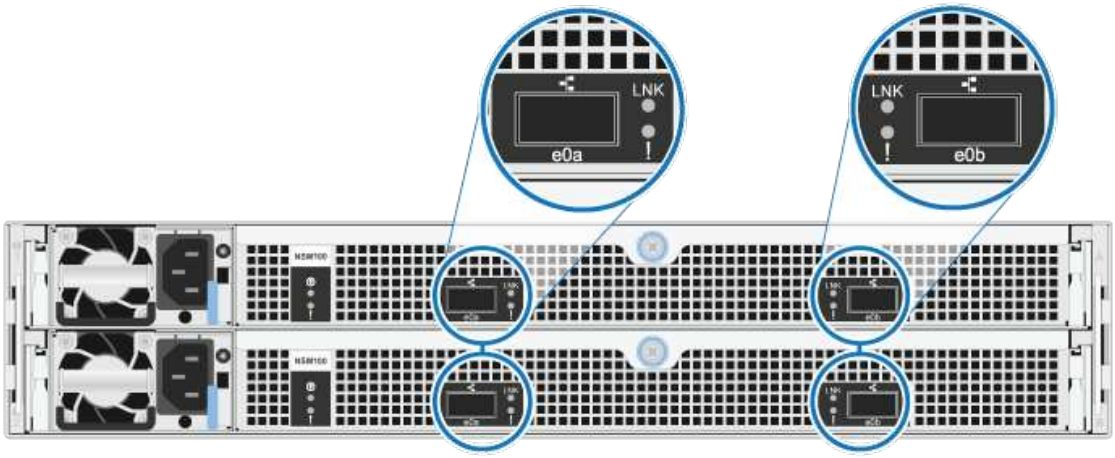

NSM100B modules

- An NS224 shelf contains two NSM100B modules. The top module goes in slot A (NSM A) and the bottom module goes in slot B (NSM B).
- Each NSM100B module includes 2 x 100GbE CX6/DX ports: e1a and e1b.

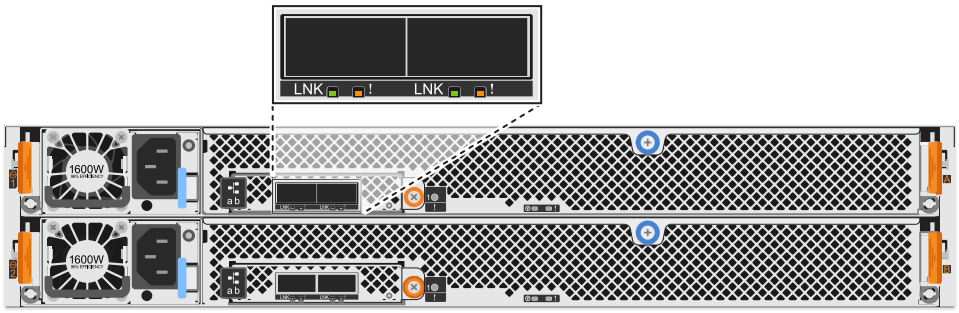

- After you have cabled a hot-added shelf, ONTAP recognizes the shelf:
  - Drive ownership is assigned if automatic drive assignment is enabled.
  - NSM shelf firmware and drive firmware should be updated automatically, if needed.

Firmware updates can take up to 30 minutes.

Cable shelf to AFF systems - NS224 shelves

You cable each NS224 shelf you are hot-adding so that each shelf has two connections to each controller in the HA pair.

About this task

Your hardware system may be compatible with both NS224 shelves with NSM100 modules and NS224 shelves with NSM100B modules. To check the compatibility and port names for your hardware and shelves, consult the NetApp Hardware Universe.

Cable shelf to AFF A1K

You can hot-add up to three additional NS224 shelves (for a total of four shelves) to an AFF A1K HA pair.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

- This procedure assumes that your HA pair has at least one existing NS224 shelf.
- This procedure addresses the following hot-add scenarios:
  - Hot-adding a second shelf to an HA pair with two RoCE-capable I/O modules in each controller. (You have installed a second I/O module and recabled the first shelf to both I/O modules or already had the first shelf cabled to two I/O modules. You will cable the second shelf to both I/O modules).
  - Hot-adding a third shelf to an HA pair with three RoCE-capable I/O modules in each controller. (You have installed a third I/O module and will cable the third shelf to only the third I/O module).
  - Hot-adding a third shelf to an HA pair with four RoCE-capable I/O modules in each controller. (You have installed a third and fourth I/O module and will cable the third shelf to the third and fourth I/O modules).
  - Hot-adding a fourth shelf to an HA pair with four RoCE-capable I/O modules in each controller. (You have installed a fourth I/O module and recabled the third shelf to the third and fourth I/O modules or already had the third shelf cabled to the third and fourth I/O modules. You will cable the fourth shelf to both the third and fourth I/O module).

Steps

1. If the NS224 shelf you are hot-adding will be the second NS224 shelf in the HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 10 port a (e10a).
   b. Cable shelf NSM A port e0b to controller B slot 11 port b (e11b).
   c. Cable shelf NSM B port e0a to controller B slot 10 port a (e10a).
   d. Cable shelf NSM B port e0b to controller A slot 11 port b (e11b).

      The following illustration highlights the cabling for the second shelf in the HA pair with two RoCE-capable I/O modules in each controller:

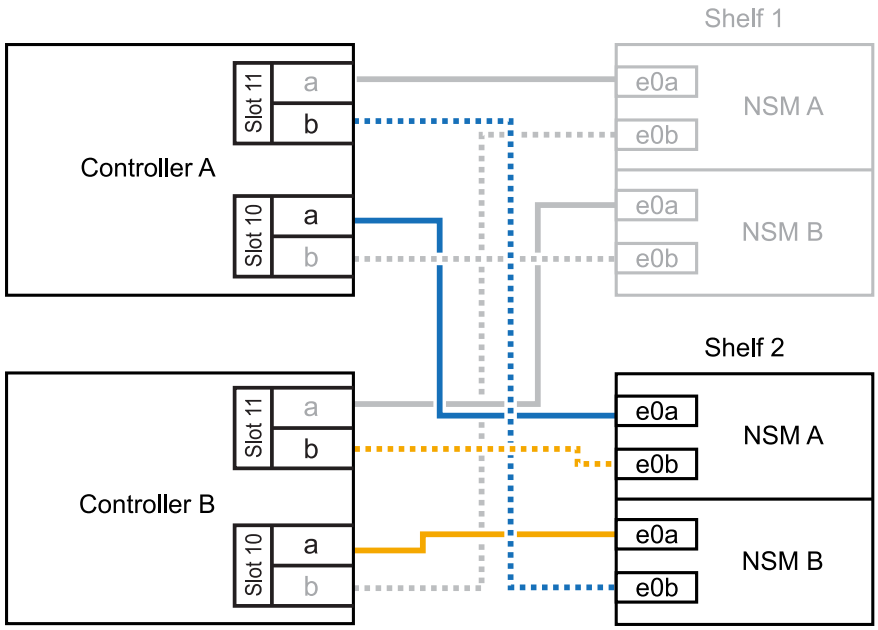

2. If the NS224 shelf you are hot-adding will be the third NS224 shelf in the HA pair with three RoCE-capable I/O modules in each controller, complete the following substeps.
   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 9 port a (e9a).

   b. Cable shelf NSM A port e0b to controller B slot 9 port b (e9b).

   c. Cable shelf NSM B port e0a to controller B slot 9 port a (e9a).

   d. Cable shelf NSM B port e0b to controller A slot 9 port b (e9b).

      The following illustration highlights the cabling for the third shelf in the HA pair with three RoCE-capable I/O modules in each controller:

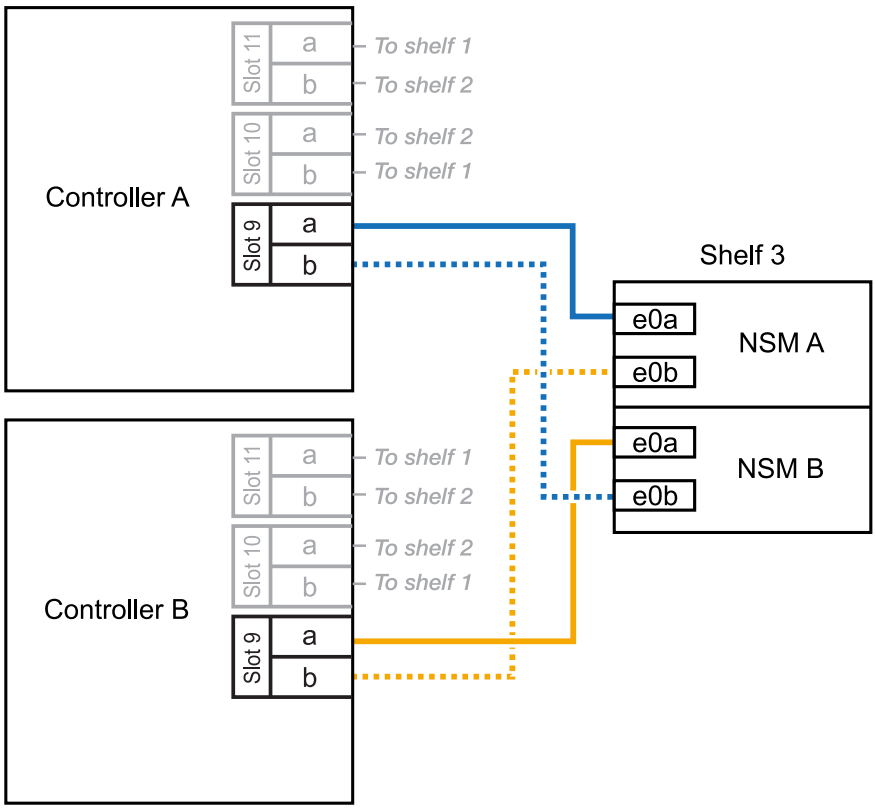

3. If the NS224 shelf you are hot-adding will be the third NS224 shelf in the HA pair with four RoCE-capable I/O modules in each controller, complete the following substeps.
   Otherwise, go to the next step.
   a. Cable shelf NSM A port e0a to controller A slot 9 port a (e9a).
   b. Cable shelf NSM A port e0b to controller B slot 8 port b (e8b).
   c. Cable shelf NSM B port e0a to controller B slot 9 port a (e9a).
   d. Cable shelf NSM B port e0b to controller A slot 8 port b (e8b).

      The following illustration highlights the cabling for the third shelf in the HA pair with four RoCE-capable I/O modules in each controller:

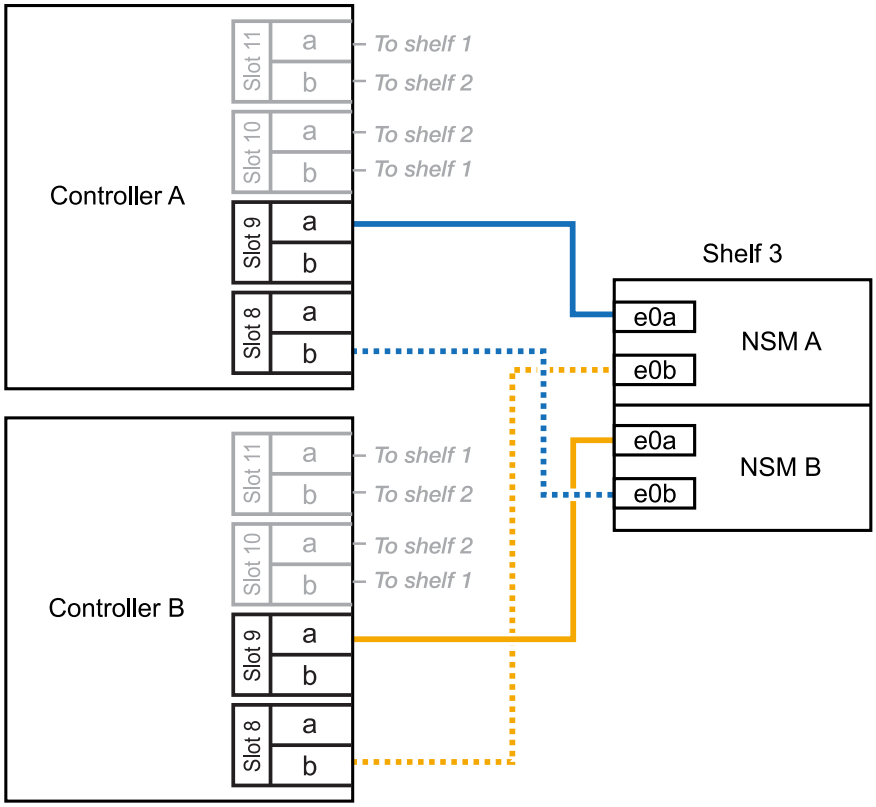

4. If the NS224 shelf you are hot-adding will be the fourth NS224 shelf in the HA pair with four RoCE-capable I/O modules in each controller, complete the following substeps.
   a. Cable shelf NSM A port e0a to controller A slot 8 port a (e8a).
   b. Cable shelf NSM A port e0b to controller B slot 9 port b (e9b).
   c. Cable shelf NSM B port e0a to controller B slot 8 port a (e8a).
   d. Cable shelf NSM B port e0b to controller A slot 9 port b (e9b).

      The following illustration highlights the cabling for the fourth shelf in the HA pair with four RoCE-capable I/O modules in each controller:

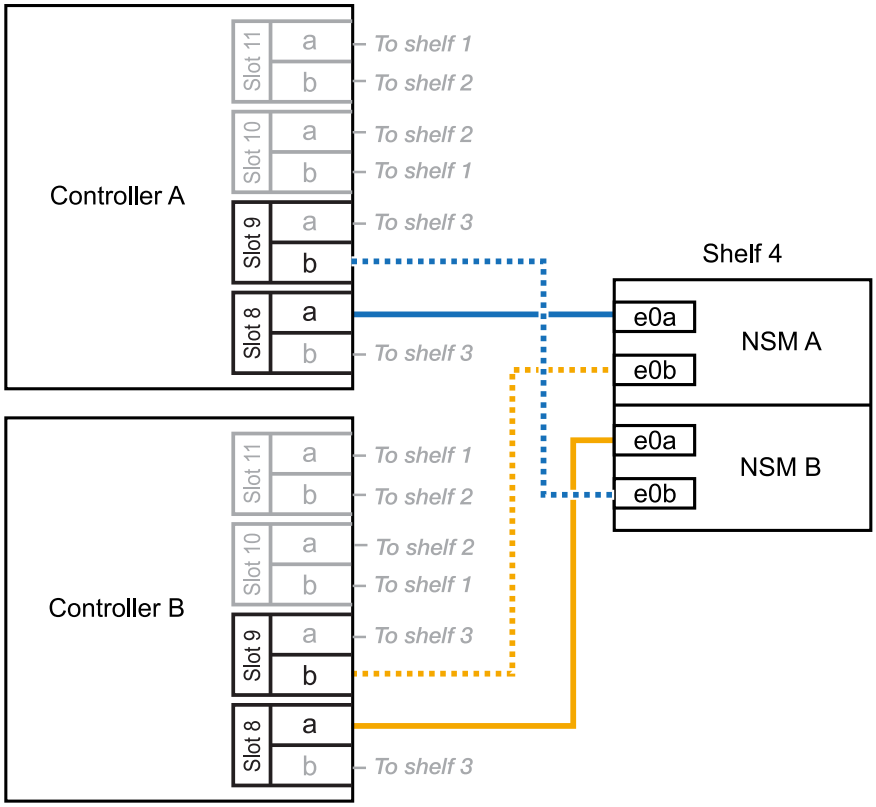

5. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to AFF A20

You can hot-add one NS224 shelf to an AFF A20 HA pair when additional storage (to the internal shelf) is needed.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

- This procedure assumes that your HA pair has only internal storage (no external shelves) and that you are hot-adding up to one additional shelf.
- This procedure addresses the following hot-add scenarios:
  - Hot-adding the first shelf to an HA pair with one RoCE-capable I/O module in each controller.
  - Hot-adding the first shelf to an HA pair with two RoCE-capable I/O modules in each controller.
- These systems are compatible with both NS224 shelves with NSM100 modules and NS224 shelves with NSM100B modules. To ensure you cable your controllers to the correct ports, replace the "X" in each diagram with the correct port number for your module:

Module type	Port labeling
NSM100	"0" ex. e0a
NSM100B	"1" ex. e1a

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (one RoCE-capable I/O module) in each controller module, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   This step assumes you installed the RoCE-capable I/O module in slot 3.

   a. Cable shelf NSM A port eXa to controller A slot 3 port a (e3a).
   b. Cable shelf NSM A port eXb to controller B slot 3 port b (e3b).
   c. Cable shelf NSM B port eXa to controller B slot 3 port a (e3a).
   d. Cable shelf NSM B port eXb to controller A slot 3 port b (e3b).

      The following illustration shows cabling for one hot-added shelf using one RoCE-capable I/O module in each controller module:

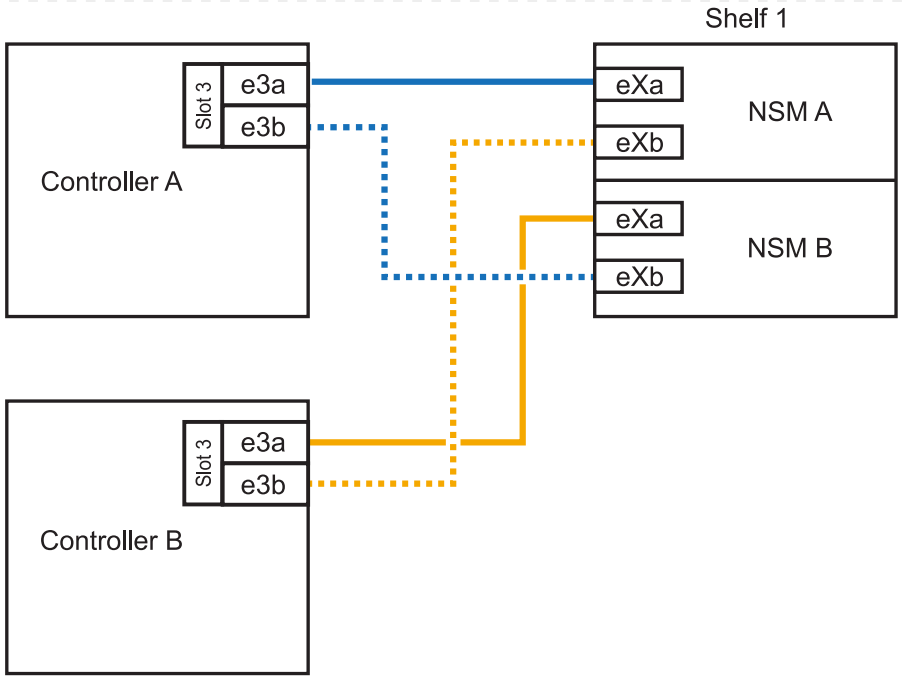

2. If you are hot-adding one shelf using two sets of RoCE-capable ports (two RoCE-capable I/O modules) in each controller module, complete the following substeps.
   a. Cable NSM A port eXa to controller A slot 3 port a (e3a).
   b. Cable NSM A port eXb to controller B slot 1 port b (e1b).
   c. Cable NSM B port eXa to controller B slot 3 port a (e3a).
   d. Cable NSM B port eXb to controller A slot 1 port b (e1b).

The following illustration shows cabling for one hot-added shelf using two RoCE-capable I/O modules in each controller module:

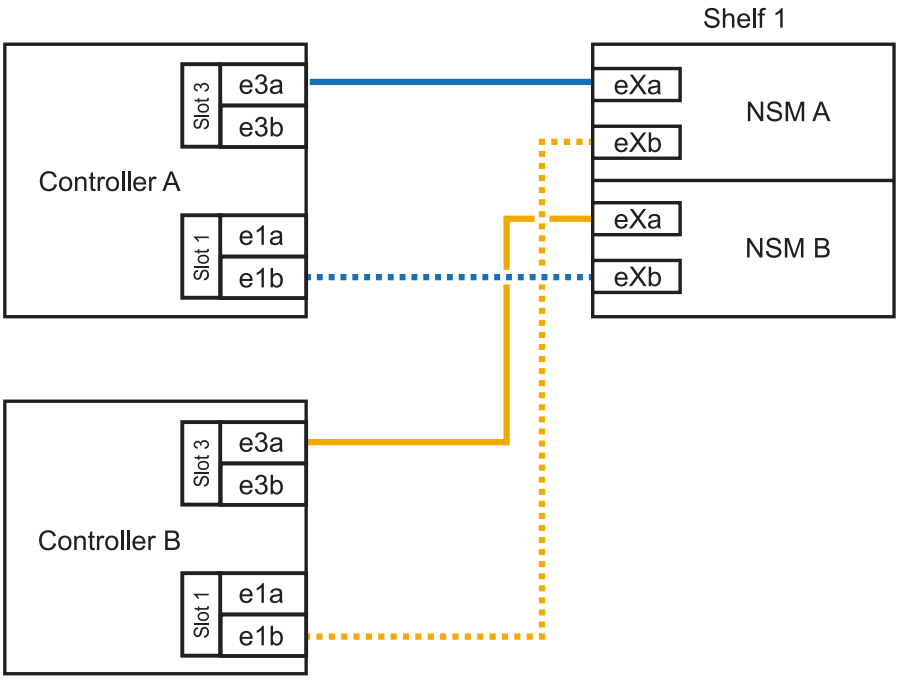

1. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to AFF A30, AFF A50, AFF C30, or AFF C60

You can hot-add up to two NS224 shelves to an AFF A30, AFF C30, AFF A50, or AFF C60 HA pair when additional storage (to the internal shelf) is needed.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

- This procedure assumes that your HA pair has only internal storage (no external shelves) and that you are Hot-adding up to two additional shelves and two RoCE-capable I/O modules in each controller.
- This procedure addresses the following hot-add scenarios:
  - Hot-adding the first shelf to an HA pair with one RoCE-capable I/O module in each controller.
  - Hot-adding the first shelf to an HA pair with two RoCE-capable I/O modules in each controller.
  - Hot- adding the second shelf to an HA pair with two RoCE-capable I/O modules in each controller.
- These systems are compatible with both NS224 shelves with NSM100 modules and NS224 shelves with NSM100B modules. To ensure you cable your controllers to the correct ports, replace the "X" in each diagram with the correct port number for your module:

Module type	Port labeling
NSM100	"0" ex. e0a
NSM100B	"1" ex. e1a

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (one RoCE-capable I/O module) in each controller module, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   This step assumes you installed the RoCE-capable I/O module in slot 3.

   a. Cable shelf NSM A port eXa to controller A slot 3 port a (e3a).
   b. Cable shelf NSM A port eXb to controller B slot 3 port b (e3b).
   c. Cable shelf NSM B port eXa to controller B slot 3 port a (e3a).
   d. Cable shelf NSM B port eXb to controller A slot 3 port b (e3b).

The following illustration shows cabling for one hot-added shelf using one RoCE-capable I/O module in each controller module:

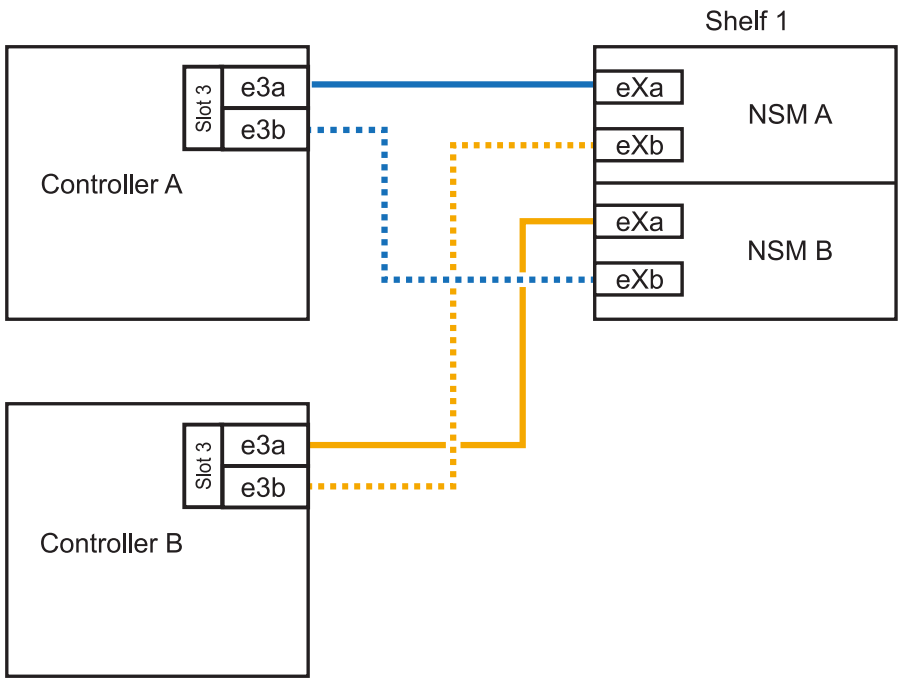

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports (two RoCE-capable I/O modules) in each controller module, complete the applicable substeps.

This step assumes you installed the RoCE-capable I/O modules in slots 3 and 1.

Shelves	Cabling
Shelf 1	a. Cable NSM A port eXa to controller A slot 3 port a (e3a). b. Cable NSM A port eXb to controller B slot 1 port b (e1b). c. Cable NSM B port eXa to controller B slot 3 port a (e3a). d. Cable NSM B port eXb to controller A slot 1 port b (e1b). e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3. The following illustration shows cabling for one hot-added shelf using two RoCE-capable I/O modules in each controller module: 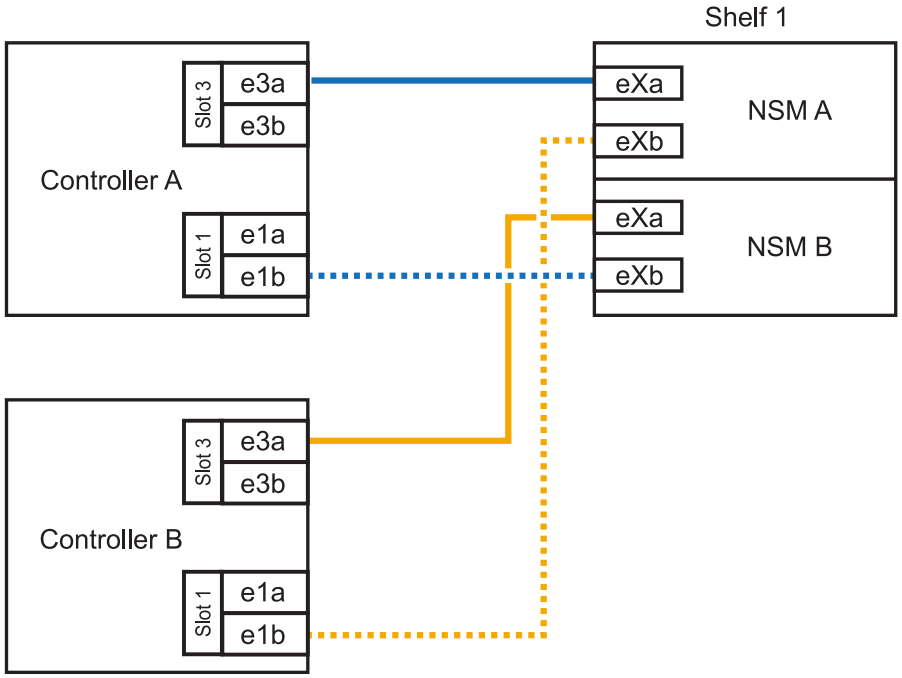

Shelves	Cabling
Shelf 2	a. Cable NSM A port eXa to controller A slot 1 port a (e1a). b. Cable NSM A port eXb to controller B slot 3 port b (e3b). c. Cable NSM B port eXa to controller B slot 1 port a (e1a). d. Cable NSM B port eXb to controller A slot 3 port b (e3b). e. Go to step 3. The following illustration shows cabling for two hot-added shelf using two RoCE-capable I/O modules in each controller module: 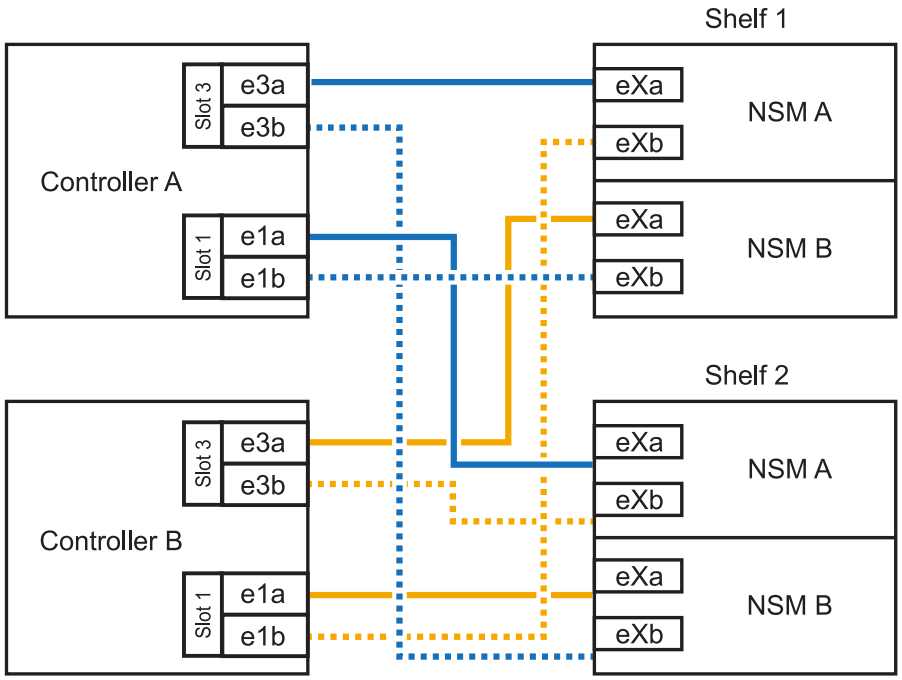

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to AFF A70, AFF A90 or AFF C80

You can hot-add up to two NS224 shelves to an AFF A70, AFF A90 or AFF C80 HA pair when additional storage (to the internal shelf) is needed.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

- This procedure assumes that your HA pair has only internal storage (no external shelves) and that you are hot-adding up to two additional shelves and two RoCE-capable I/O modules in each controller.
- This procedure addresses the following hot-add scenarios:
  - Hot-adding the first shelf to an HA pair with one RoCE-capable I/O module in each controller.
  - Hot-adding the first shelf to an HA pair with two RoCE-capable I/O modules in each controller.
  - Hot- adding the second shelf to an HA pair with two RoCE-capable I/O modules in each controller.

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (one RoCE-capable I/O module) in each controller module, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   This step assumes you installed the RoCE-capable I/O module in slot 11.

   a. Cable shelf NSM A port e0a to controller A slot 11 port a (e11a).
   b. Cable shelf NSM A port e0b to controller B slot 11 port b (e11b).
   c. Cable shelf NSM B port e0a to controller B slot 11 port a (e11a).
   d. Cable shelf NSM B port e0b to controller A slot 11 port b (e11b).

      The following illustration shows cabling for one hot-added shelf using one RoCE-capable I/O module in each controller module:

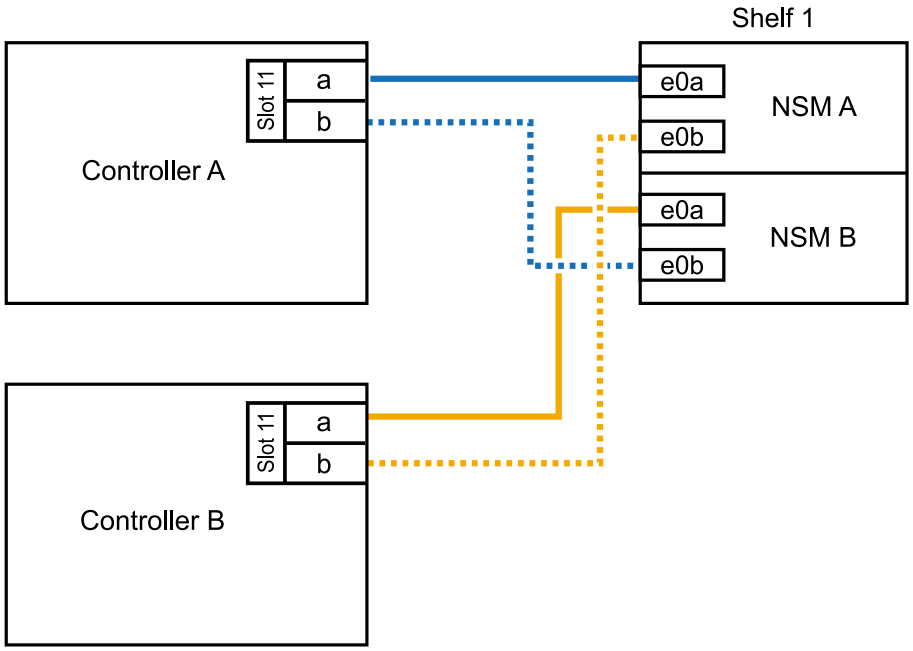

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports (two RoCE-capable I/O modules) in each controller module, complete the applicable substeps.

This step assumes you installed the RoCE-capable I/O modules in slots 11 and 8.

Shelves	Cabling
Shelf 1	a. Cable NSM A port e0a to controller A slot 11 port a (e11a). b. Cable NSM A port e0b to controller B slot 8 port b (e8b). c. Cable NSM B port e0a to controller B slot 11 port a (e11a). d. Cable NSM B port e0b to controller A slot 8 port b (e8b). e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3. The following illustration shows cabling for one hot-added shelf using two RoCE-capable I/O modules in each controller module: 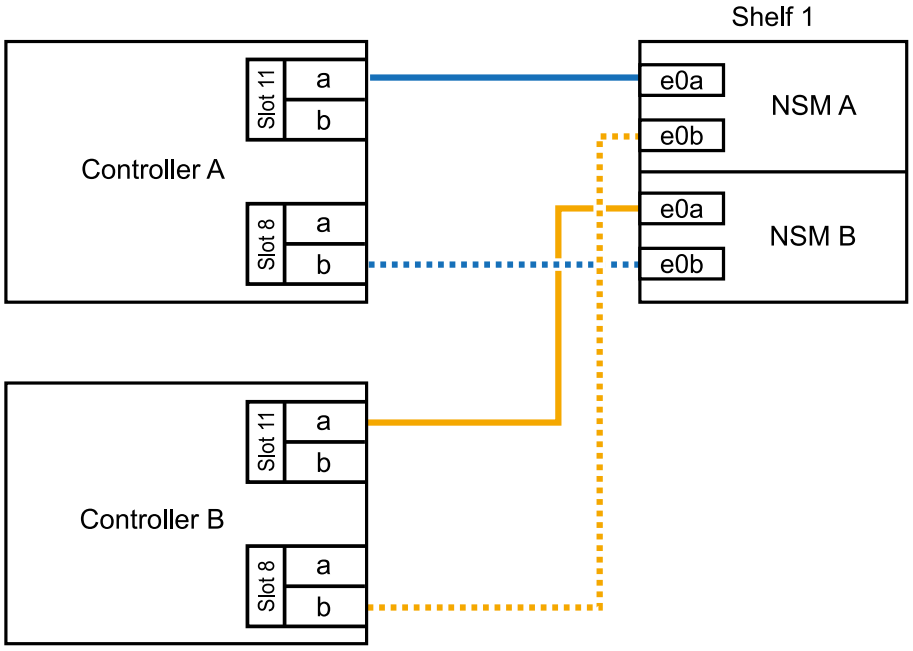

Shelves	Cabling
Shelf 2	a. Cable NSM A port e0a to controller A slot 8 port a (e8a). b. Cable NSM A port e0b to controller B slot 11 port b (e11b). c. Cable NSM B port e0a to controller B slot 8 port a (e8a). d. Cable NSM B port e0b to controller A slot 11 port b (e11b). e. Go to step 3. The following illustration shows cabling for two hot-added shelf using two RoCE-capable I/O modules in each controller module: 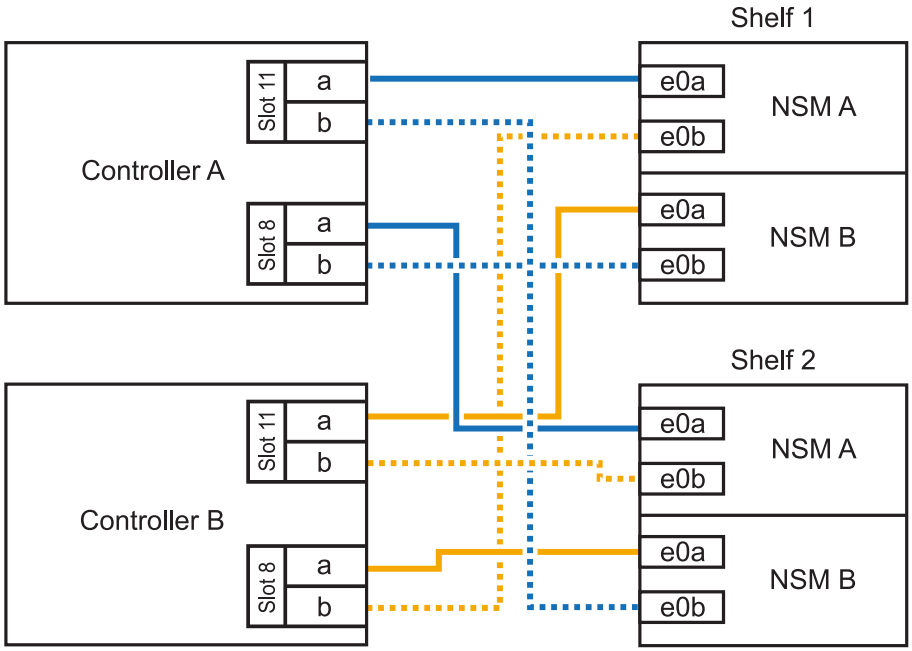

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to AFF A250 or AFF C250

When additional storage is needed, you can hot-add a maximum of one NS224 shelf to an AFF A250 or AFF C250 HA pair.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

When viewed from the rear of the platform chassis, the RoCE-capable card port on the left is port "a" (e1a) and the port on the right is port "b" (e1b).

Steps

1. Cable the shelf connections:
   a. Cable shelf NSM A port e0a to controller A slot 1 port a (e1a).
   b. Cable shelf NSM A port e0b to controller B slot 1 port b (e1b).
   c. Cable shelf NSM B port e0a to controller B slot 1 port a (e1a).
   d. Cable shelf NSM B port e0b to controller A slot 1 port b (e1b).

   The following illustration shows the shelf cabling when completed.

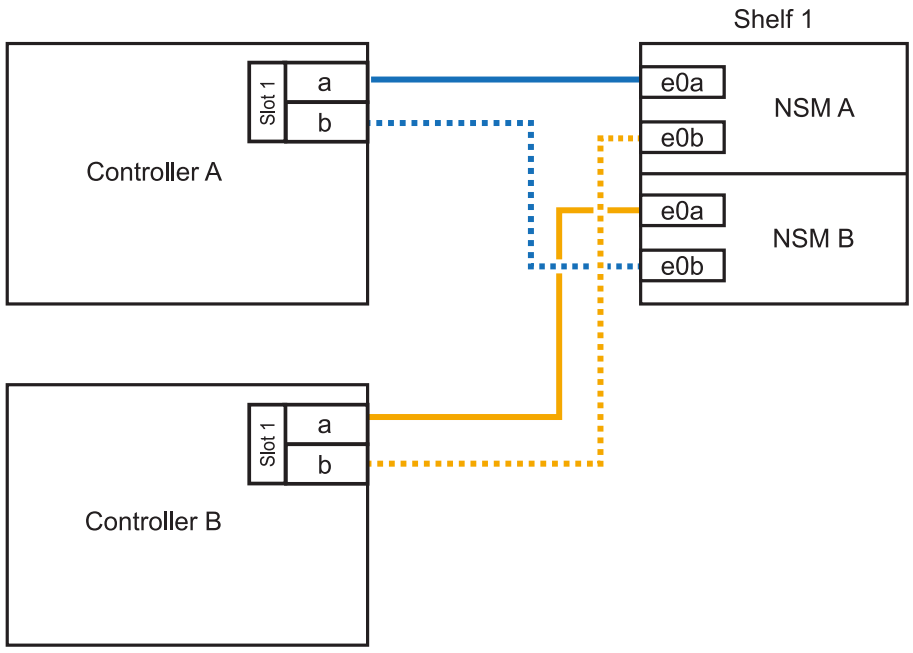

2. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to

Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to AFF A400 or AFF C400

How you cable an NS224 shelf for a hot-add depends on whether you have an AFF A400 or AFF C400 HA pair.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

Cable shelf to an AFF A400 HA pair

For an AFF A400 HA pair, you can hot-add up to two shelves and use onboard ports e0c/e0d and ports in slot 5 as needed.

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (onboard RoCE-capable ports) on each controller, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A port e0c.

   b. Cable shelf NSM A port e0b to controller B port e0d.

   c. Cable shelf NSM B port e0a to controller B port e0c.

   d. Cable shelf NSM B port e0b to controller A port e0d.

      The following illustration shows cabling for one hot-added shelf using one set of RoCE-capable ports on each controller:

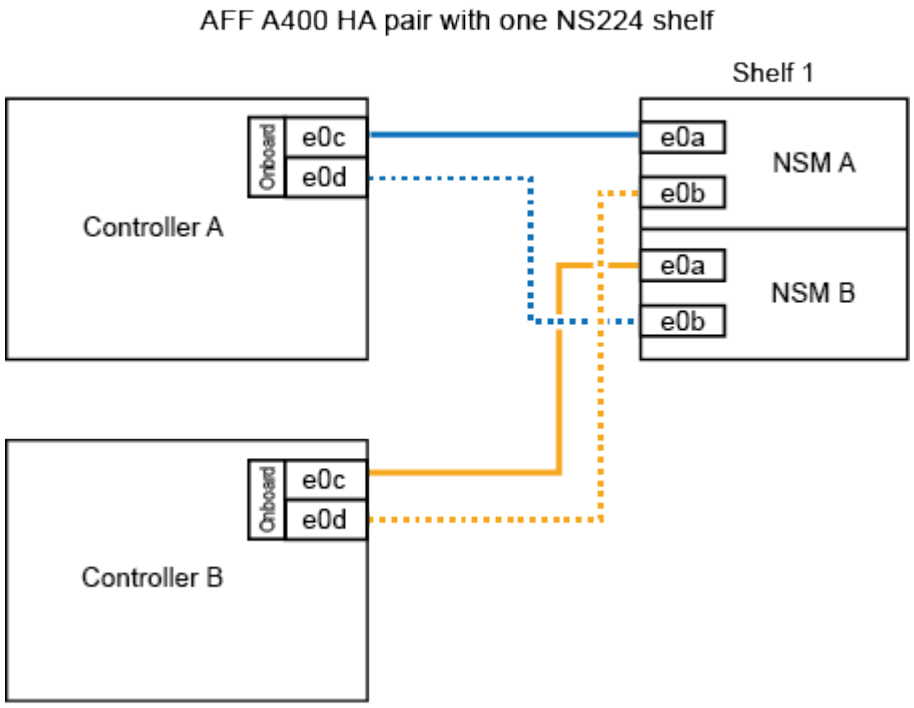

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports (on board and PCIe card RoCE-capable ports) on each controller, complete the following substeps.

Shelves	Cabling
Shelf 1	a. Cable NSM A port e0a to controller A port e0c. b. Cable NSM A port e0b to controller B slot 5 port 2 (e5b). c. Cable NSM B port e0a to controller B port e0c. d. Cable NSM B port e0b to controller A slot 5 port 2 (e5b). e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3.
Shelf 2	a. Cable NSM A port e0a to controller A slot 5 port 1 (e5a). b. Cable NSM A port e0b to controller B port e0d. c. Cable NSM B port e0a to controller B slot 5 port 1 (e5a). d. Cable NSM B port e0b to controller A port e0d. e. Go to step 3.

The following illustration shows cabling for two hot-added shelves:

AFF A400 HA pair with two NS224 shelves

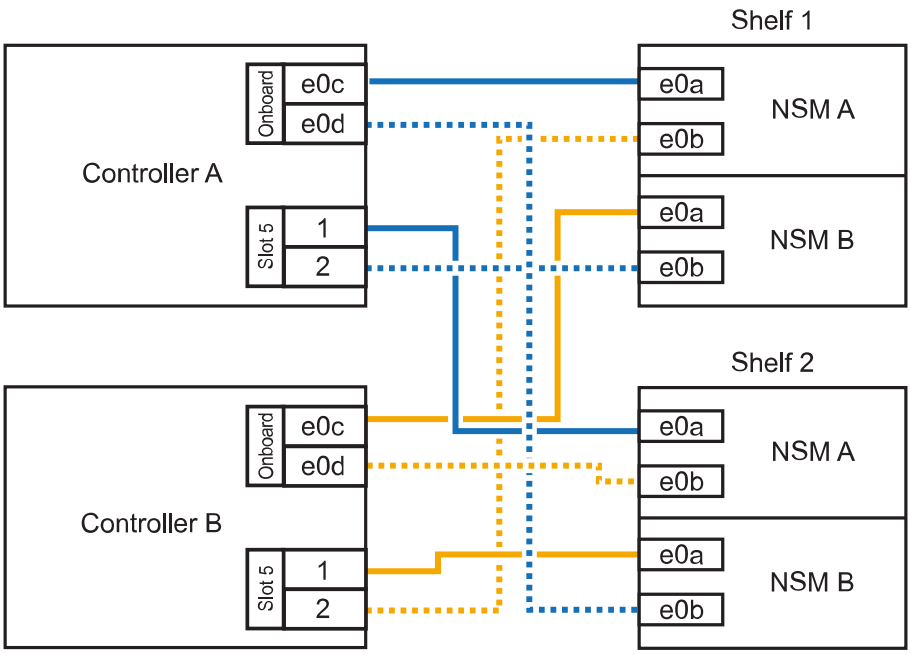

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

4. If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then re enable automatic drive assignment, if needed. See Complete the hot-add.

Otherwise, you are done with this procedure.

Cable shelf to an AFF C400 HA pair

For an AFF C400 HA pair, you can hot-add up to two shelves and use ports in slot 4 and 5 as needed.

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports on each controller, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 4 port 1 (e4a).
   b. Cable shelf NSM A port e0b to controller B slot 4 port 2 (e4b).
   c. Cable shelf NSM B port e0a to controller B slot 4 port 1 (e4a).
   d. Cable shelf NSM B port e0b to controller A slot 4 port 2 (e4b).

      The following illustration shows cabling for one hot-added shelf using one set of RoCE-capable ports on each controller:

      AFF C400 HA pair with one NS224 shelf

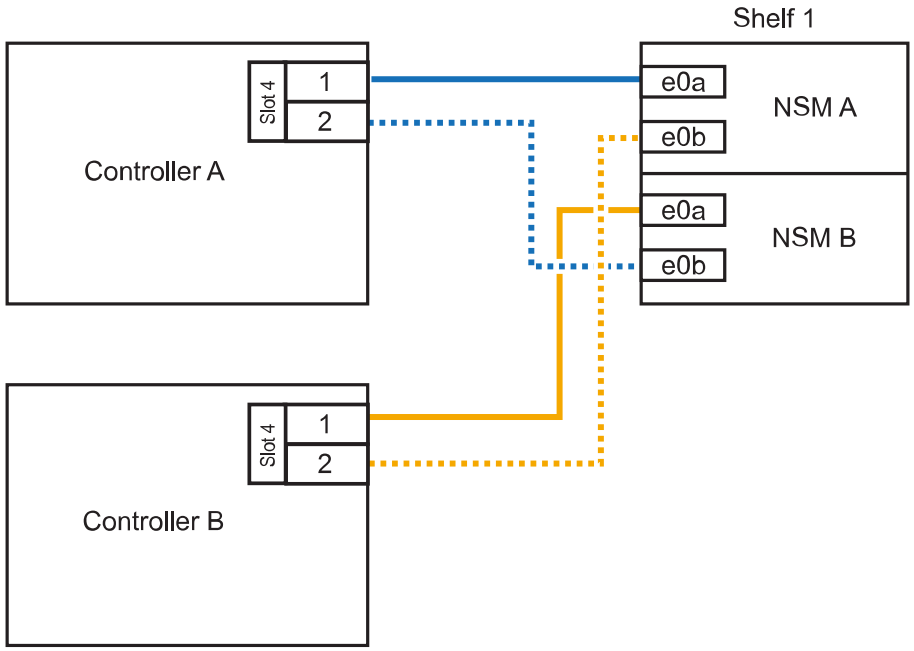

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports on each controller, complete the following substeps.

Shelves	Cabling
Shelf 1	a. Cable NSM A port e0a to controller A slot 4 port 1 (e4a). b. Cable NSM A port e0b to controller B slot 5 port 2 (e5b). c. Cable NSM B port e0a to controller B port slot 4 port 1 (e4a). d. Cable NSM B port e0b to controller A slot 5 port 2 (e5b). e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3.
Shelf 2	a. Cable NSM A port e0a to controller A slot 5 port 1 (e5a). b. Cable NSM A port e0b to controller B slot 4 port 2 (e4b). c. Cable NSM B port e0a to controller B slot 5 port 1 (e5a). d. Cable NSM B port e0b to controller A slot 4 port 2 (e4b). e. Go to step 3.

The following illustration shows cabling for two hot-added shelves:

AFF C400 HA pair with two NS224 shelves

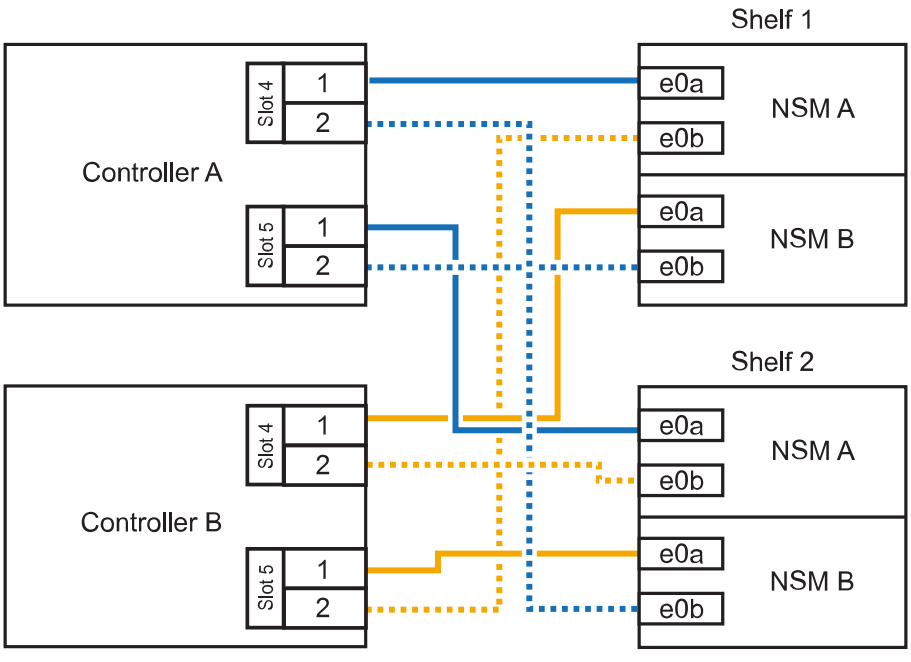

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What’s next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to AFF A800 or AFF C800

How you cable an NS224 shelf in an AFF A800 or AFF C800 HA pair depends on the number of shelves you are hot-adding and the number of RoCE-capable port sets (one or two) you are using on the controllers.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (one RoCE-capable PCIe card) on each controller, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   This step assumes you installed the RoCE-capable PCIe card in slot 5.

   a. Cable shelf NSM A port e0a to controller A slot 5 port a (e5a).
   b. Cable shelf NSM A port e0b to controller B slot 5 port b (e5b).
   c. Cable shelf NSM B port e0a to controller B slot 5 port a (e5a).
   d. Cable shelf NSM B port e0b to controller A slot 5 port b (e5b).

   The following illustration shows cabling for one hot-added shelf using one RoCE-capable PCIe card on each controller:

   AFF A800 or AFF C800 HA pair with one NS224 shelf

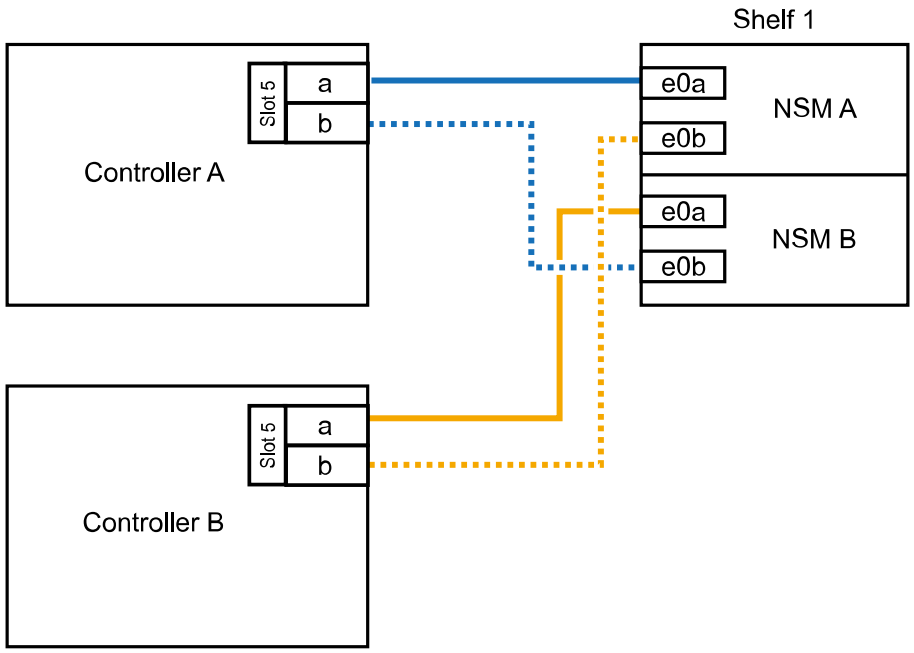

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports (two RoCE-capable

PCIe cards) on each controller, complete the applicable substeps.

This step assumes you installed the RoCE-capable PCIe cards in slot 5 and slot 3.

Shelves	Cabling
Shelf 1	These substeps assume you are beginning the cabling by cabling shelf port e0a to the RoCE-capable PCIe card in slot 5, instead of slot 3. a. Cable NSM A port e0a to controller A slot 5 port a (e5a). b. Cable NSM A port e0b to controller B slot 3 port b (e3b). c. Cable NSM B port e0a to controller B slot 5 port a (e5a). d. Cable NSM B port e0b to controller A slot 3 port b (e3b). e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3.
Shelf 2	These substeps assume you are beginning the cabling by cabling shelf port e0a to the RoCE-capable PCIe card in slot 3, instead of slot 5 (which correlates with the cabling substeps for shelf 1). a. Cable NSM A port e0a to controller A slot 3 port a (e3a). b. Cable NSM A port e0b to controller B slot 5 port b (e5b). c. Cable NSM B port e0a to controller B slot 3 port a (e3a). d. Cable NSM B port e0b to controller A slot 5 port b (e5b). e. Go to step 3.

The following illustration shows cabling for two hot-added shelves:

AFF A800 or AFF C800 HA pair with two NS224 shelves

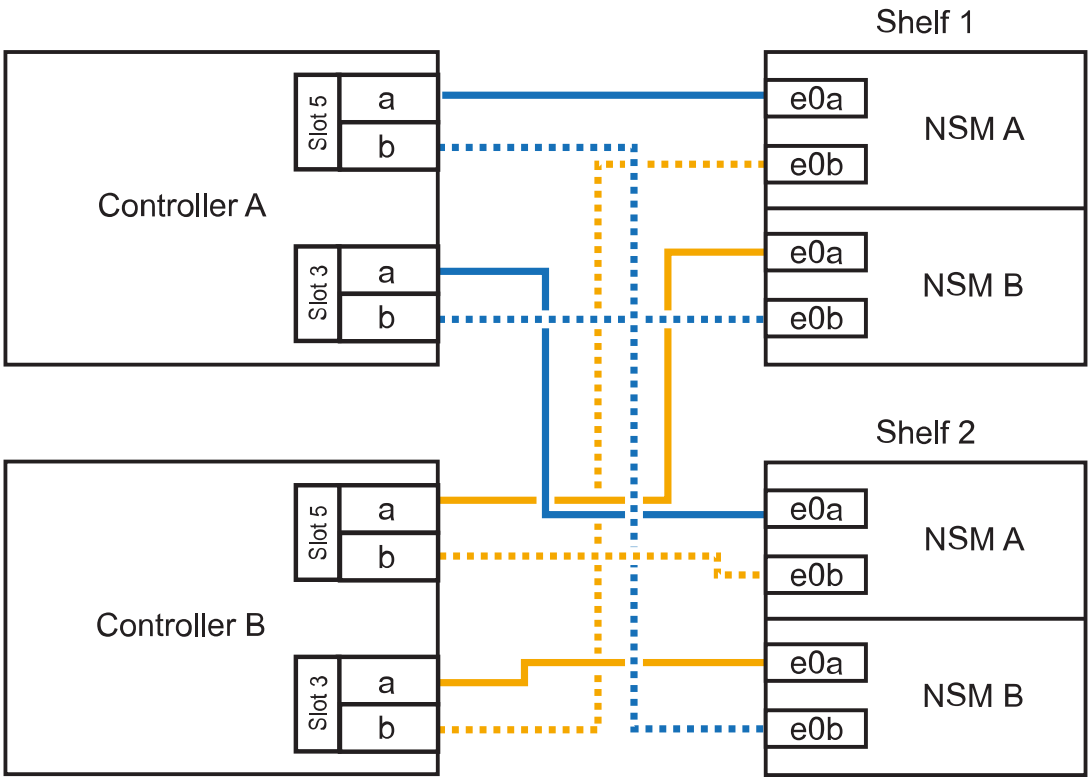

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to AFF A900

When additional storage is needed, you can hot-add up to three additional NS224 drive shelves (for a total of four shelves) to an AFF A900 HA pair.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

- This procedure assumes that your HA pair has at least one existing NS224 shelf and that you are hot-adding up to three additional shelves.
- If your HA pair has only one existing NS224 shelf, this procedure assumes that the shelf is cabled across two RoCE-capable 100GbE I/O modules on each controller.

Steps

1. If the NS224 shelf you are hot-adding will be the second NS224 shelf in the HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 10 port a (e10a).
   b. Cable shelf NSM A port e0b to controller B slot 2 port b (e2b).
   c. Cable shelf NSM B port e0a to controller B slot 10 port a (e10a).
   d. Cable shelf NSM B port e0b to controller A slot 2 port b (e2b).

      The following illustration shows the second shelf cabling (and the first shelf).

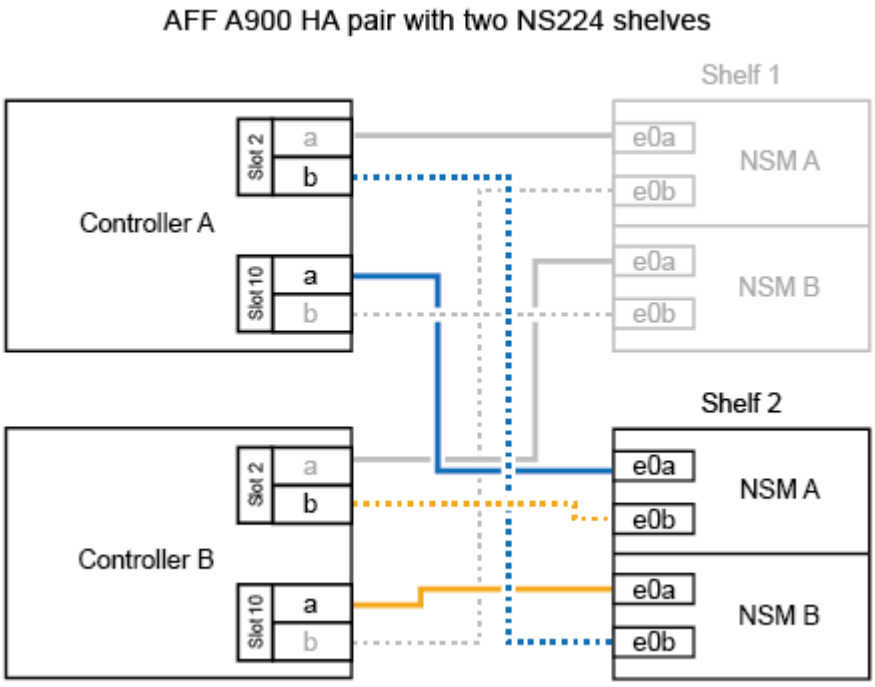

2. If the NS224 shelf you are hot-adding will be the third NS224 shelf in the HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 1 port a (e1a).
   b. Cable shelf NSM A port e0b to controller B slot 11 port b (e11b).
   c. Cable shelf NSM B port e0a to controller B slot 1 port a (e1a).
   d. Cable shelf NSM B port e0b to controller A slot 11 port b (e11b).

      The following illustration shows the third shelf cabling.

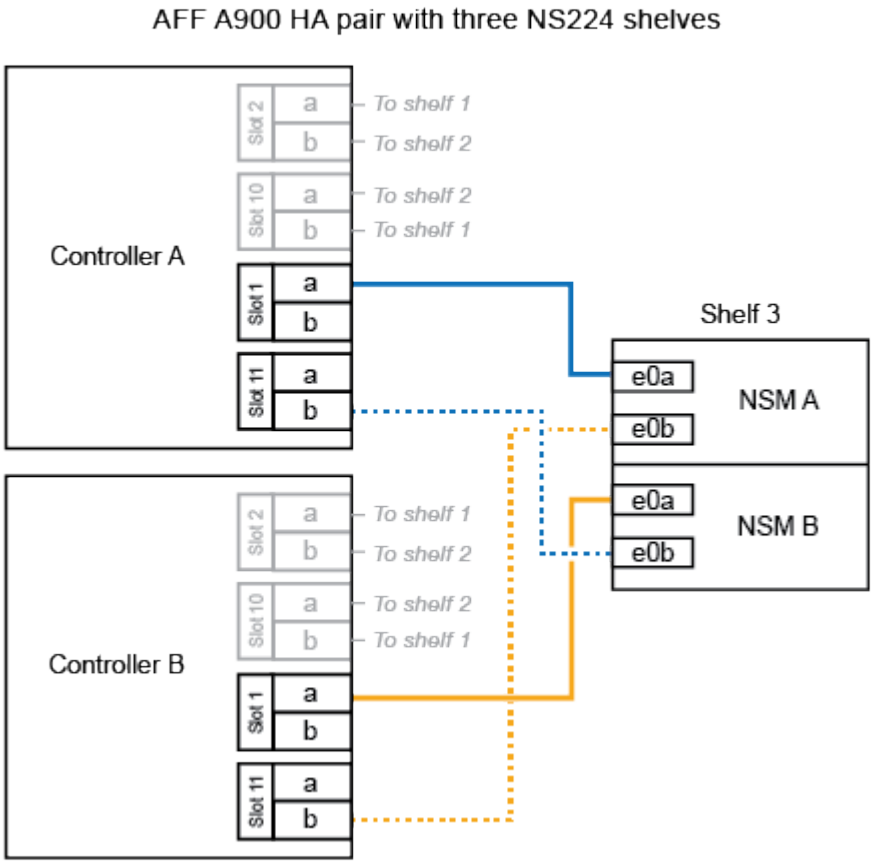

3. If the NS224 shelf you are hot-adding will be the fourth NS224 shelf in the HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 11 port a (e11a).
   b. Cable shelf NSM A port e0b to controller B slot 1 port b (e1b).
   c. Cable shelf NSM B port e0a to controller B slot 11 port a (e11a).
   d. Cable shelf NSM B port e0b to controller A slot 1 port b (e1b).

      The following illustration shows the fourth shelf cabling.

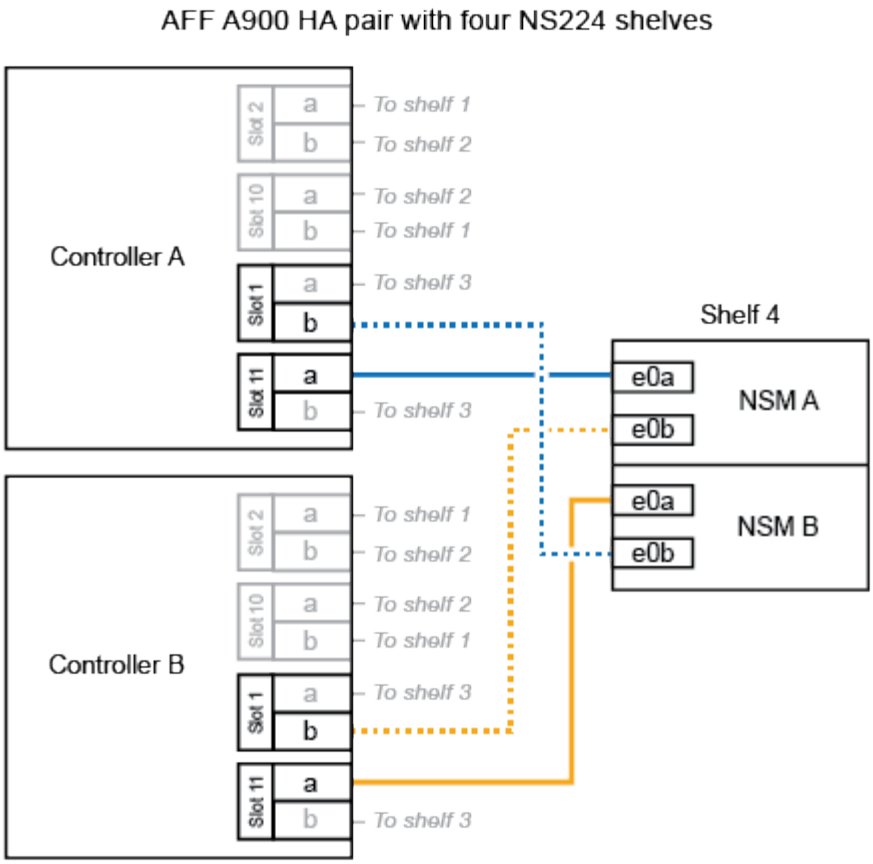

4. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to ASA systems - NS224 shelves

You cable each NS224 shelf you are hot-adding so that each shelf has two connections to each controller in the HA pair.

About this task

Your hardware system may be compatible with both NS224 shelves with NSM100 modules and NS224 shelves with NSM100B modules. To check the compatibility and port names for your hardware and shelves, consult the NetApp Hardware Universe.

Cable shelf to ASA A1K

You can hot-add up to three additional NS224 shelves (for a total of four shelves) to an ASA A1K HA pair.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

- This procedure assumes that your HA pair has at least one existing NS224 shelf.
- This procedure addresses the following hot-add scenarios:
  - Hot-adding a second shelf to an HA pair with two RoCE-capable I/O modules in each controller. (You have installed a second I/O module and recabled the first shelf to both I/O modules or already had the first shelf cabled to two I/O modules. You will cable the second shelf to both I/O modules).
  - Hot-adding a third shelf to an HA pair with three RoCE-capable I/O modules in each controller. (You have installed a third I/O module and will cable the third shelf to only the third I/O module).
  - Hot-adding a third shelf to an HA pair with four RoCE-capable I/O modules in each controller. (You have installed a third and fourth I/O module and will cable the third shelf to the third and fourth I/O modules).
  - Hot-adding a fourth shelf to an HA pair with four RoCE-capable I/O modules in each controller. (You have installed a fourth I/O module and recabled the third shelf to the third and fourth I/O modules or already had the third shelf cabled to the third and fourth I/O modules. You will cable the fourth shelf to both the third and fourth I/O module).

Steps

1. If the NS224 shelf you are hot-adding will be the second NS224 shelf in the HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 10 port a (e10a).
   b. Cable shelf NSM A port e0b to controller B slot 11 port b (e11b).
   c. Cable shelf NSM B port e0a to controller B slot 10 port a (e10a).
   d. Cable shelf NSM B port e0b to controller A slot 11 port b (e11b).

      The following illustration highlights the cabling for the second shelf in the HA pair with two RoCE-capable I/O modules in each controller:

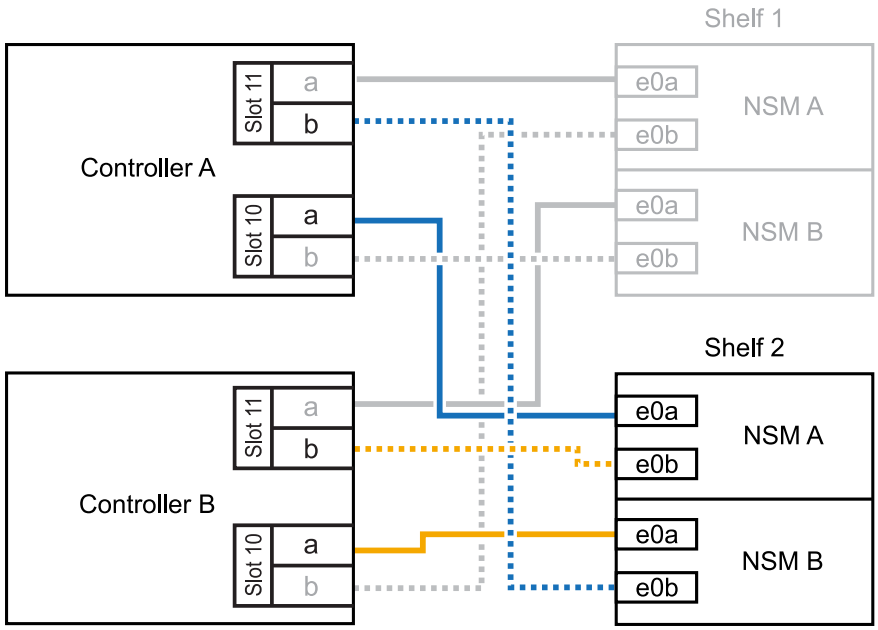

2. If the NS224 shelf you are hot-adding will be the third NS224 shelf in the HA pair with three RoCE-capable I/O modules in each controller, complete the following substeps.
   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 9 port a (e9a).

   b. Cable shelf NSM A port e0b to controller B slot 9 port b (e9b).

   c. Cable shelf NSM B port e0a to controller B slot 9 port a (e9a).

   d. Cable shelf NSM B port e0b to controller A slot 9 port b (e9b).

      The following illustration highlights the cabling for the third shelf in the HA pair with three RoCE-capable I/O modules in each controller:

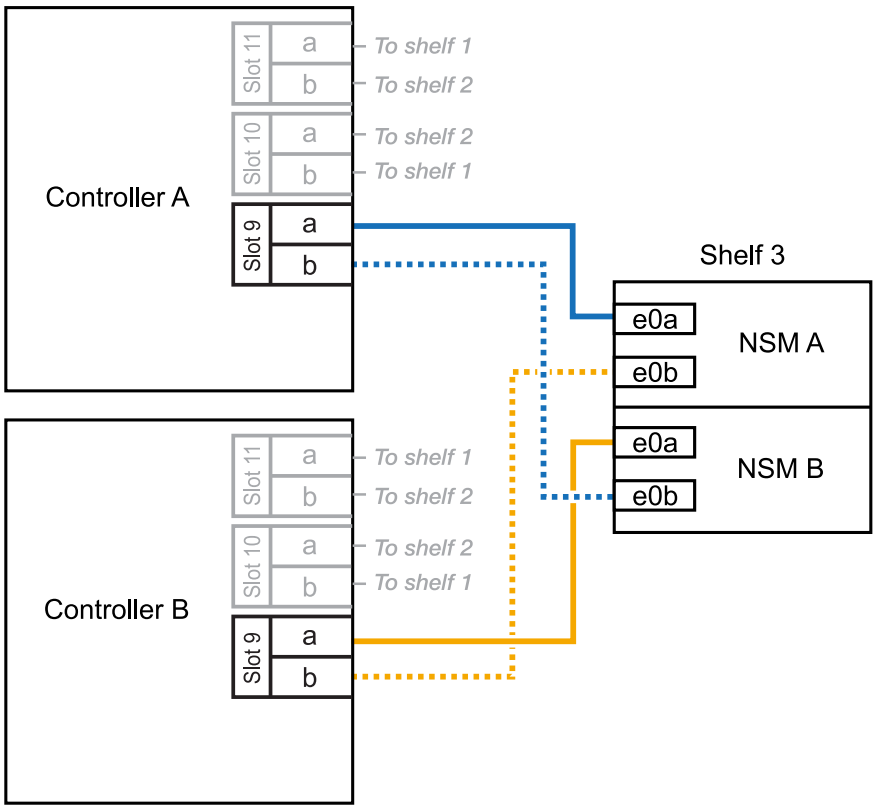

3. If the NS224 shelf you are hot-adding will be the third NS224 shelf in the HA pair with four RoCE-capable I/O modules in each controller, complete the following substeps.
   Otherwise, go to the next step.
   a. Cable shelf NSM A port e0a to controller A slot 9 port a (e9a).
   b. Cable shelf NSM A port e0b to controller B slot 8 port b (e8b).
   c. Cable shelf NSM B port e0a to controller B slot 9 port a (e9a).
   d. Cable shelf NSM B port e0b to controller A slot 8 port b (e8b).

      The following illustration highlights the cabling for the third shelf in the HA pair with four RoCE-capable I/O modules in each controller:

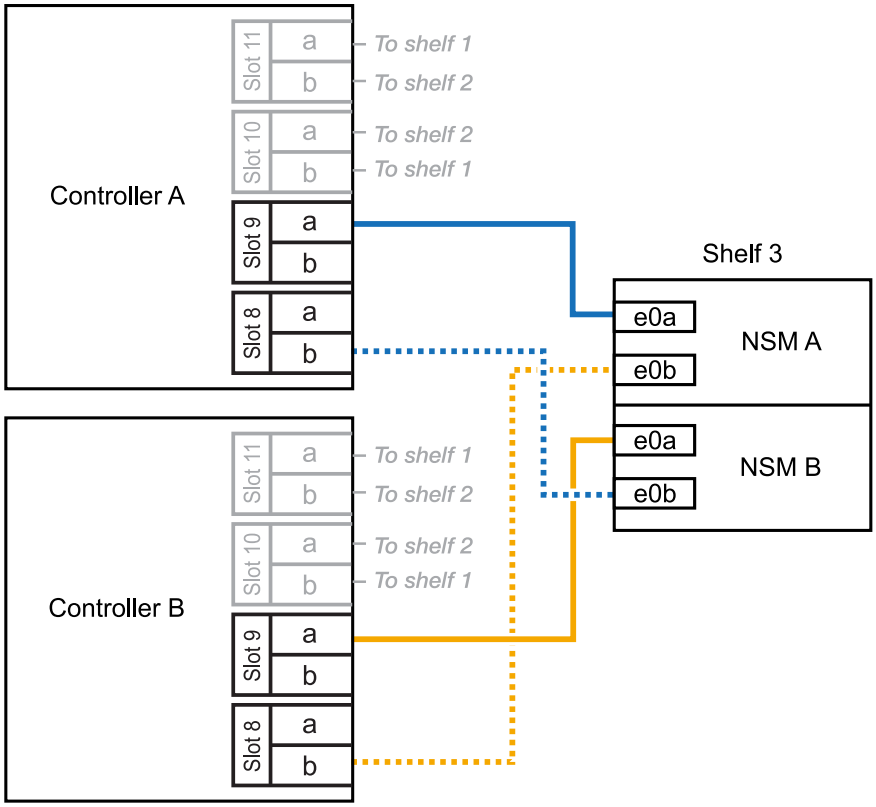

4. If the NS224 shelf you are hot-adding will be the fourth NS224 shelf in the HA pair with four RoCE-capable I/O modules in each controller, complete the following substeps.
   a. Cable shelf NSM A port e0a to controller A slot 8 port a (e8a).
   b. Cable shelf NSM A port e0b to controller B slot 9 port b (e9b).
   c. Cable shelf NSM B port e0a to controller B slot 8 port a (e8a).
   d. Cable shelf NSM B port e0b to controller A slot 9 port b (e9b).

      The following illustration highlights the cabling for the fourth shelf in the HA pair with four RoCE-capable I/O modules in each controller:

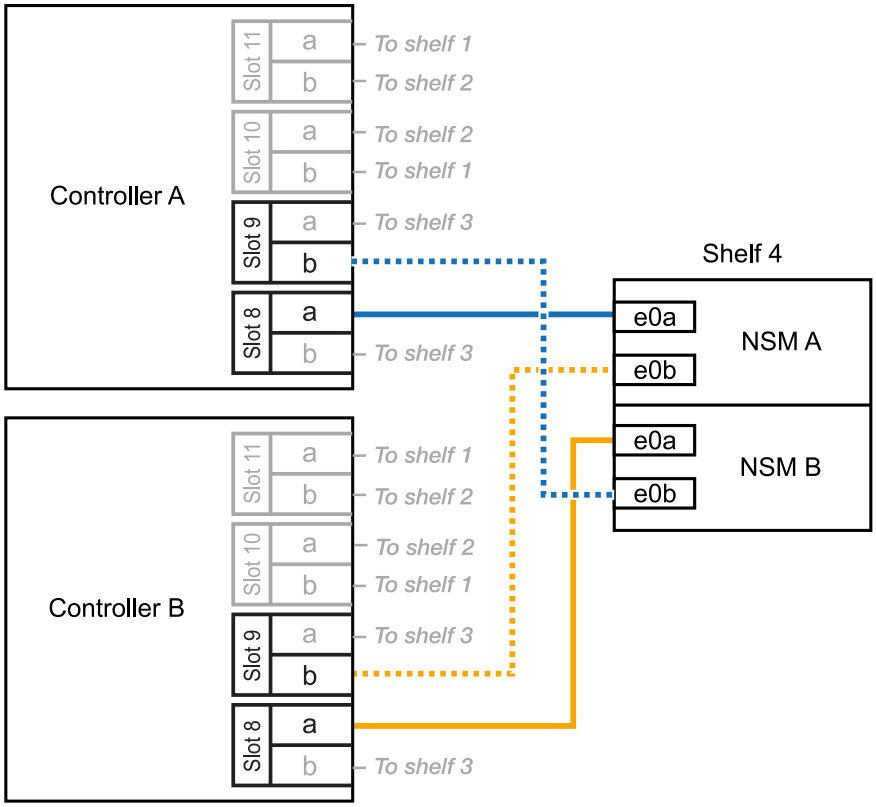

5. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to ASA A20

You can hot-add a maximum of one NS224 shelf to an ASA A20 HA pair when additional storage (to the internal shelf) is needed.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

- This procedure assumes that your HA pair has only internal storage (no external shelves) and that you are hot-adding up to one additional shelf.
- This procedure addresses the following hot-add scenarios:
  - Hot-adding the first shelf to an HA pair with one RoCE-capable I/O module in each controller.
  - Hot-adding the first shelf to an HA pair with two RoCE-capable I/O modules in each controller.
- These systems are compatible with both NS224 shelves with NSM100 modules and NS224 shelves with NSM100B modules. To ensure you cable your controllers to the correct ports, replace the "X" in each diagram with the correct port number for your module:

Module type	Port labeling
NSM100	"0" ex. e0a
NSM100B	"1" ex. e1a

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (one RoCE-capable I/O module) in each controller module, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   This step assumes you installed the RoCE-capable I/O module in slot 3.

   a. Cable shelf NSM A port eXa to controller A slot 3 port a (e3a).
   b. Cable shelf NSM A port eXb to controller B slot 3 port b (e3b).
   c. Cable shelf NSM B port eXa to controller B slot 3 port a (e3a).
   d. Cable shelf NSM B port eXb to controller A slot 3 port b (e3b).

      The following illustration shows cabling for one hot-added shelf using one RoCE-capable I/O module in each controller module:

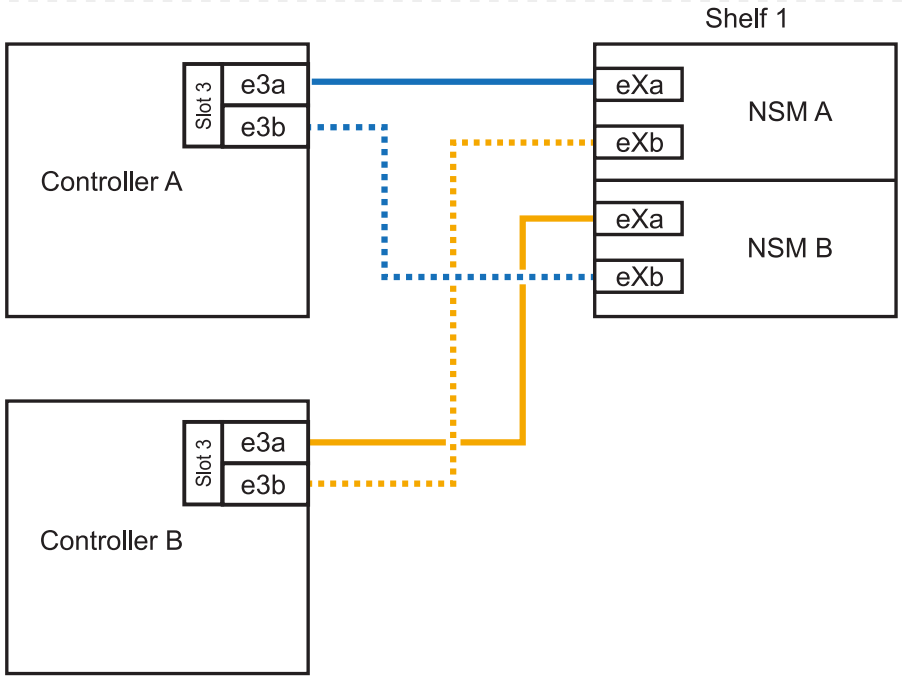

2. If you are hot-adding one shelf using two sets of RoCE-capable ports (two RoCE-capable I/O modules) in each controller module, complete the following substeps.
   a. Cable NSM A port eXa to controller A slot 3 port a (e3a).
   b. Cable NSM A port eXb to controller B slot 1 port b (e1b).
   c. Cable NSM B port eXa to controller B slot 3 port a (e3a).
   d. Cable NSM B port eXb to controller A slot 1 port b (e1b).

The following illustration shows cabling for one hot-added shelf using two RoCE-capable I/O modules in each controller module:

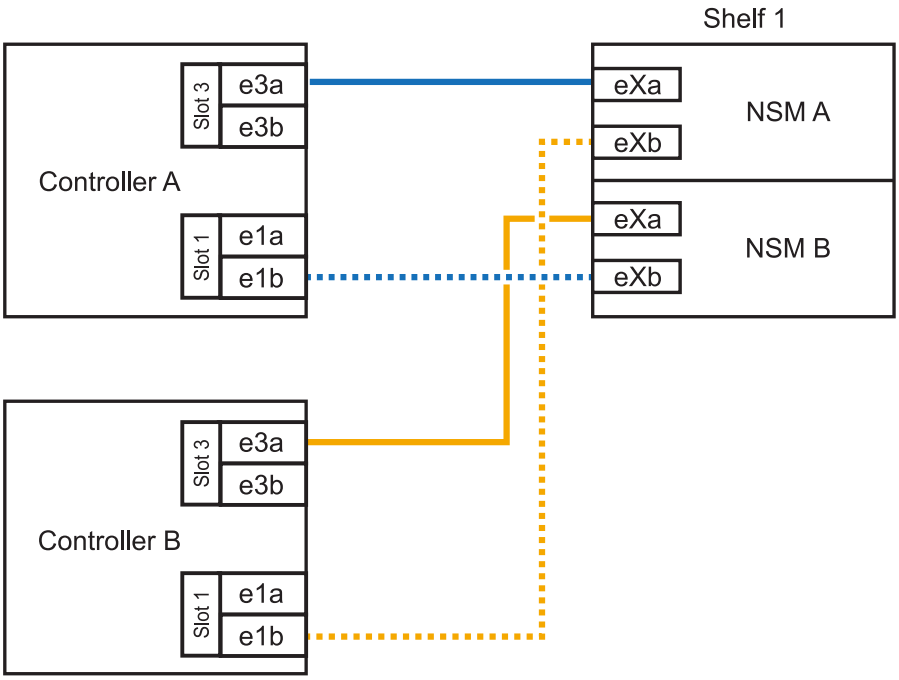

1. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to ASA A30 or ASA A50

You can hot-add up to two NS224 shelves to an ASA A30 or A50 HA pair when additional storage (to the internal shelf) is needed.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

- This procedure assumes that your HA pair has only internal storage (no external shelves) and that you are Hot-adding up to two additional shelves and two RoCE-capable I/O modules in each controller.
- This procedure addresses the following hot-add scenarios:
  - Hot-adding the first shelf to an HA pair with one RoCE-capable I/O module in each controller.
  - Hot-adding the first shelf to an HA pair with two RoCE-capable I/O modules in each controller.
  - Hot- adding the second shelf to an HA pair with two RoCE-capable I/O modules in each controller.
- These systems are compatible with both NS224 shelves with NSM100 modules and NS224 shelves with NSM100B modules. To ensure you cable your controllers to the correct ports, replace the "X" in each diagram with the correct port number for your module:

Module type	Port labeling
NSM100	"0" ex. e0a
NSM100B	"1" ex. e1a

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (one RoCE-capable I/O module) in each controller module, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   This step assumes you installed the RoCE-capable I/O module in slot 3.

   a. Cable shelf NSM A port eXa to controller A slot 3 port a (e3a).
   b. Cable shelf NSM A port eXb to controller B slot 3 port b (e3b).
   c. Cable shelf NSM B port eXa to controller B slot 3 port a (e3a).
   d. Cable shelf NSM B port eXb to controller A slot 3 port b (e3b).

The following illustration shows cabling for one hot-added shelf using one RoCE-capable I/O module in each controller module:

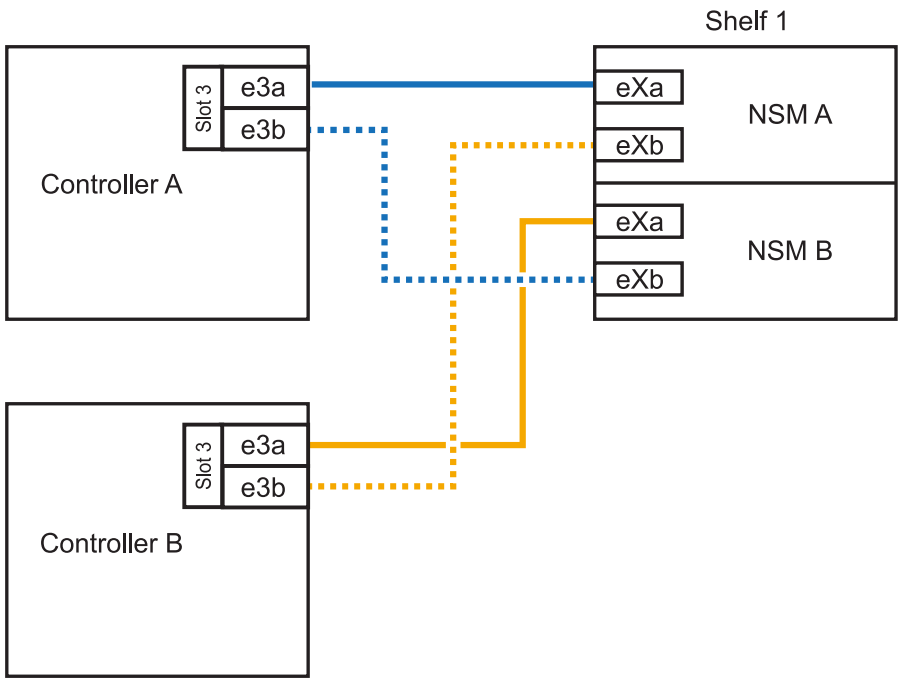

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports (two RoCE-capable I/O modules) in each controller module, complete the applicable substeps.

This step assumes you installed the RoCE-capable I/O modules in slots 3 and 1.

Shelves	Cabling
Shelf 1	a. Cable NSM A port eXa to controller A slot 3 port a (e3a). b. Cable NSM A port eXb to controller B slot 1 port b (e1b). c. Cable NSM B port eXa to controller B slot 3 port a (e3a). d. Cable NSM B port eXb to controller A slot 1 port b (e1b). e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3. The following illustration shows cabling for one hot-added shelf using two RoCE-capable I/O modules in each controller module: 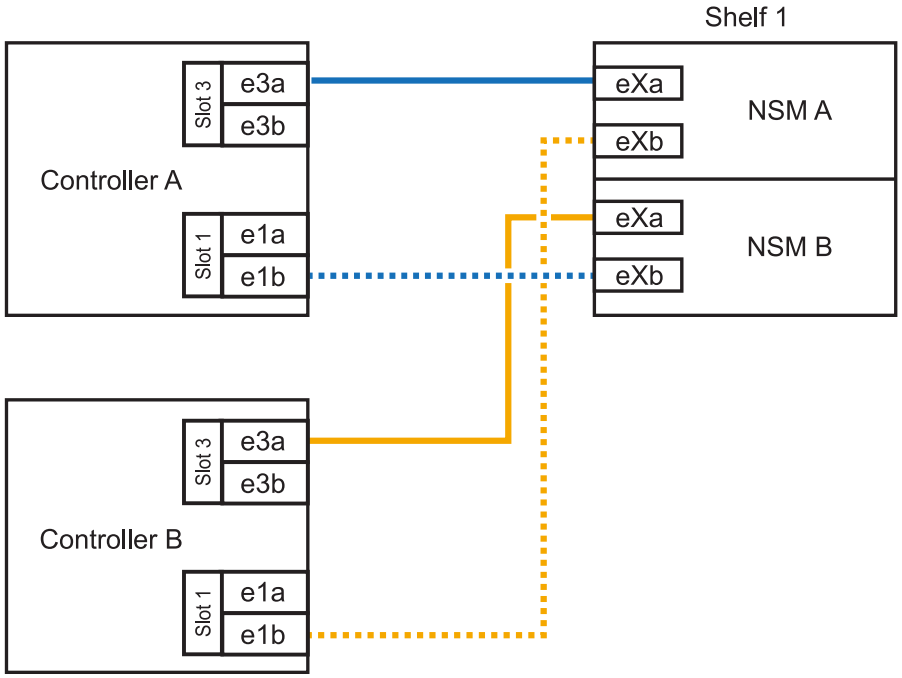

Shelves	Cabling
Shelf 2	a. Cable NSM A port eXa to controller A slot 1 port a (e1a). b. Cable NSM A port eXb to controller B slot 3 port b (e3b). c. Cable NSM B port eXa to controller B slot 1 port a (e1a). d. Cable NSM B port eXb to controller A slot 3 port b (e3b). e. Go to step 3. The following illustration shows cabling for two hot-added shelf using two RoCE-capable I/O modules in each controller module: 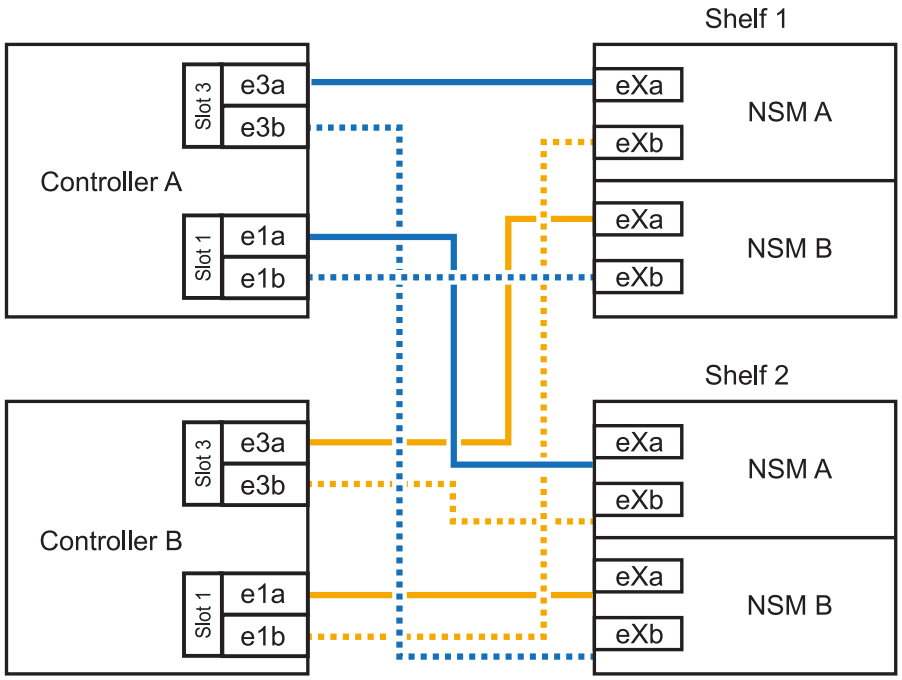

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to ASA A70 or ASA A90

You can hot-add up to two NS224 shelves to an ASA A70 or ASA A90 HA pair when additional storage (to the internal shelf) is needed.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

- This procedure assumes that your HA pair has only internal storage (no external shelves) and that you are hot-adding up to two additional shelves and two RoCE-capable I/O modules in each controller.
- This procedure addresses the following hot-add scenarios:
  - Hot-adding the first shelf to an HA pair with one RoCE-capable I/O module in each controller.
  - Hot-adding the first shelf to an HA pair with two RoCE-capable I/O modules in each controller.
  - Hot- adding the second shelf to an HA pair with two RoCE-capable I/O modules in each controller.

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (one RoCE-capable I/O module) in each controller module, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   This step assumes you installed the RoCE-capable I/O module in slot 11.

   a. Cable shelf NSM A port e0a to controller A slot 11 port a (e11a).
   b. Cable shelf NSM A port e0b to controller B slot 11 port b (e11b).
   c. Cable shelf NSM B port e0a to controller B slot 11 port a (e11a).
   d. Cable shelf NSM B port e0b to controller A slot 11 port b (e11b).

      The following illustration shows cabling for one hot-added shelf using one RoCE-capable I/O module in each controller module:

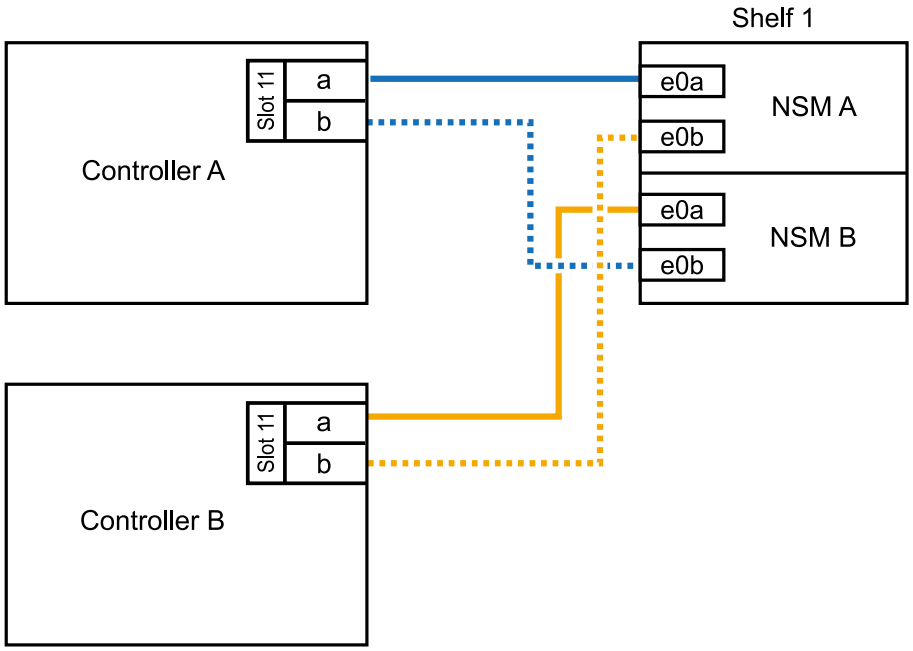

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports (two RoCE-capable I/O modules) in each controller module, complete the applicable substeps.

This step assumes you installed the RoCE-capable I/O modules in slots 11 and 8.

Shelves	Cabling
Shelf 1	a. Cable NSM A port e0a to controller A slot 11 port a (e11a). b. Cable NSM A port e0b to controller B slot 8 port b (e8b). c. Cable NSM B port e0a to controller B slot 11 port a (e11a). d. Cable NSM B port e0b to controller A slot 8 port b (e8b). e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3. The following illustration shows cabling for one hot-added shelf using two RoCE-capable I/O modules in each controller module: 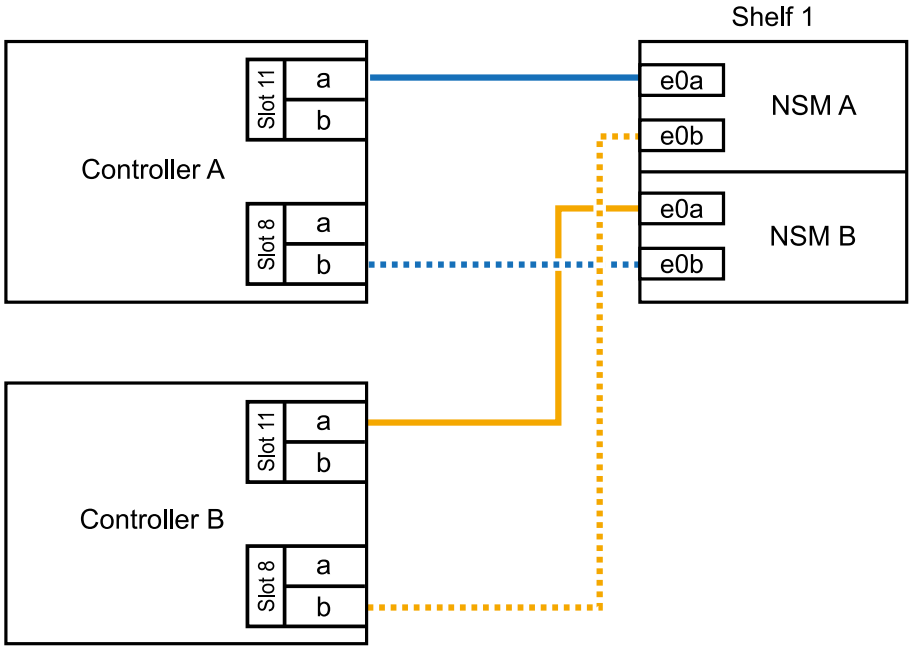

Shelves	Cabling
Shelf 2	a. Cable NSM A port e0a to controller A slot 8 port a (e8a). b. Cable NSM A port e0b to controller B slot 11 port b (e11b). c. Cable NSM B port e0a to controller B slot 8 port a (e8a). d. Cable NSM B port e0b to controller A slot 11 port b (e11b). e. Go to step 3. The following illustration shows cabling for two hot-added shelf using two RoCE-capable I/O modules in each controller module: 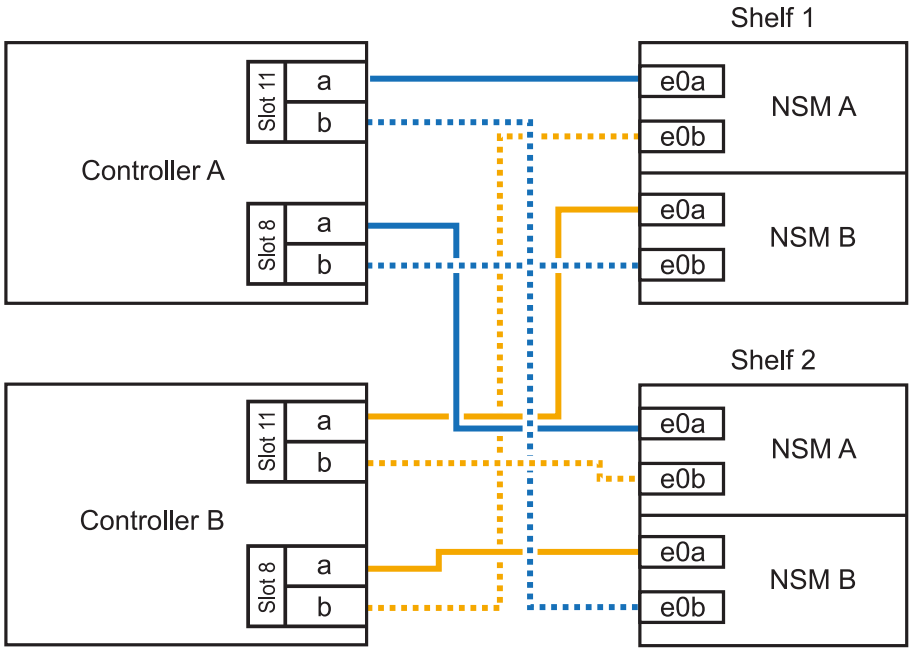

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to ASA A250 or ASA C250

When additional storage is needed, you can hot-add a maximum of one NS224 shelf to an ASA A250 or ASA C250 HA pair.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

When viewed from the rear of the platform chassis, the RoCE-capable card port on the left is port "a" (e1a) and the port on the right is port "b" (e1b).

Steps

1. Cable the shelf connections:
   a. Cable shelf NSM A port e0a to controller A slot 1 port a (e1a).
   b. Cable shelf NSM A port e0b to controller B slot 1 port b (e1b).
   c. Cable shelf NSM B port e0a to controller B slot 1 port a (e1a).
   d. Cable shelf NSM B port e0b to controller A slot 1 port b (e1b).

   The following illustration shows the shelf cabling when completed.

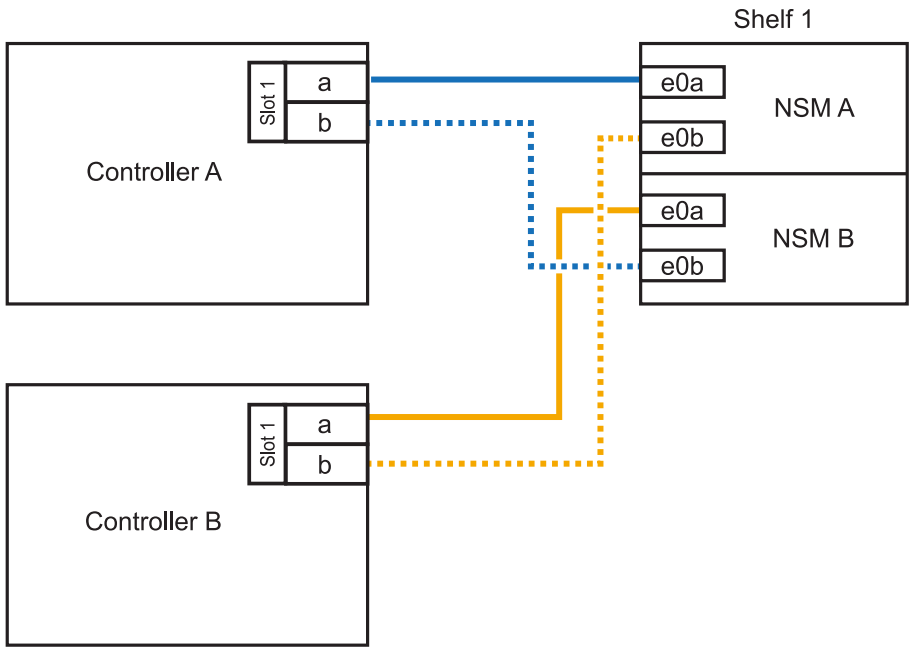

2. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to

Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to ASA A400 or ASA C400

How you cable an NS224 shelf for a hot-add depends on whether you have an ASA A400 or ASA C400 HA pair.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

Cable shelf to an AFF A400 HA pair

For an AFF A400 HA pair, you can hot-add up to two shelves and use onboard ports e0c/e0d and ports in slot 5 as needed.

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (onboard RoCE-capable ports) on each controller, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A port e0c.
   b. Cable shelf NSM A port e0b to controller B port e0d.
   c. Cable shelf NSM B port e0a to controller B port e0c.
   d. Cable shelf NSM B port e0b to controller A port e0d.

      The following illustration shows cabling for one hot-added shelf using one set of RoCE-capable ports on each controller:

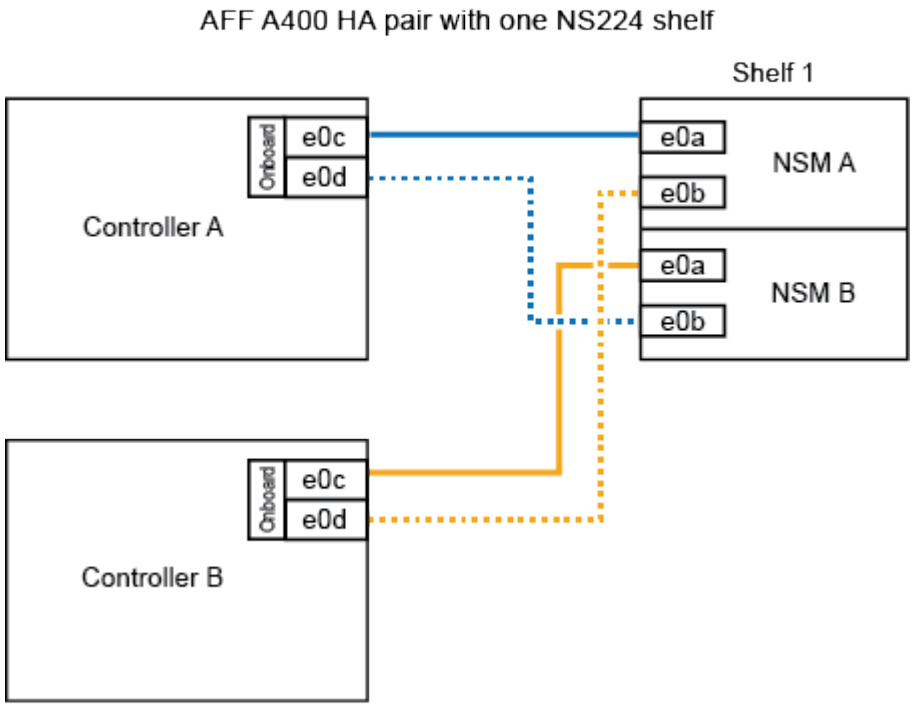

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports (on board and PCIe card RoCE-capable ports) on each controller, complete the following substeps.

Shelves	Cabling
Shelf 1	a. Cable NSM A port e0a to controller A port e0c. b. Cable NSM A port e0b to controller B slot 5 port 2 (e5b). c. Cable NSM B port e0a to controller B port e0c. d. Cable NSM B port e0b to controller A slot 5 port 2 (e5b). e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3.
Shelf 2	a. Cable NSM A port e0a to controller A slot 5 port 1 (e5a). b. Cable NSM A port e0b to controller B port e0d. c. Cable NSM B port e0a to controller B slot 5 port 1 (e5a). d. Cable NSM B port e0b to controller A port e0d. e. Go to step 3.

The following illustration shows cabling for two hot-added shelves:

AFF A400 HA pair with two NS224 shelves

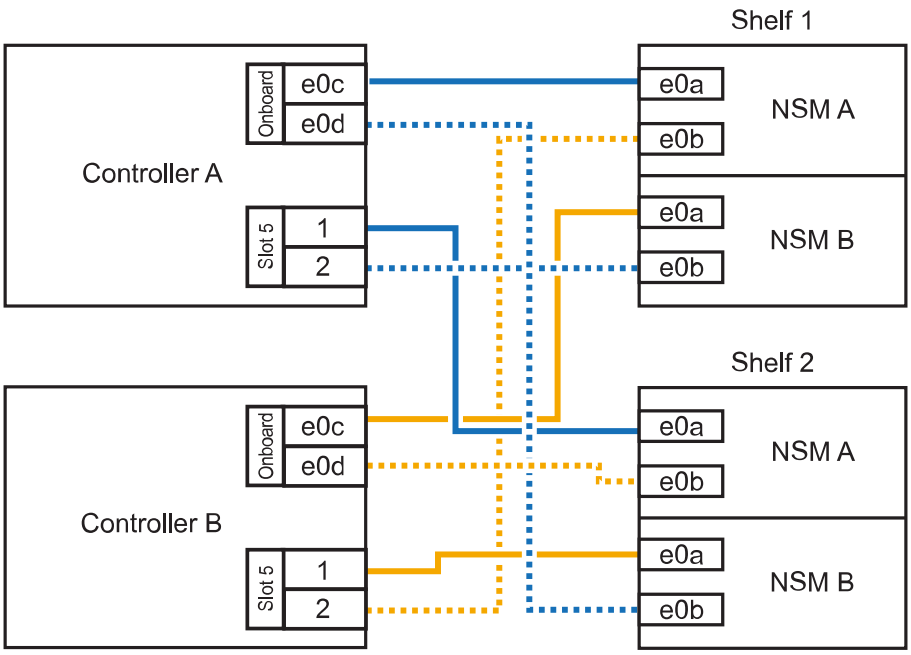

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

4. If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then re enable automatic drive assignment, if needed. See Complete the hot-add.

Otherwise, you are done with this procedure.

Cable shelf to an AFF C400 HA pair

For an AFF C400 HA pair, you can hot-add up to two shelves and use ports in slot 4 and 5 as needed.

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports on each controller, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 4 port 1 (e4a).
   b. Cable shelf NSM A port e0b to controller B slot 4 port 2 (e4b).
   c. Cable shelf NSM B port e0a to controller B slot 4 port 1 (e4a).
   d. Cable shelf NSM B port e0b to controller A slot 4 port 2 (e4b).

      The following illustration shows cabling for one hot-added shelf using one set of RoCE-capable ports on each controller:

      AFF C400 HA pair with one NS224 shelf

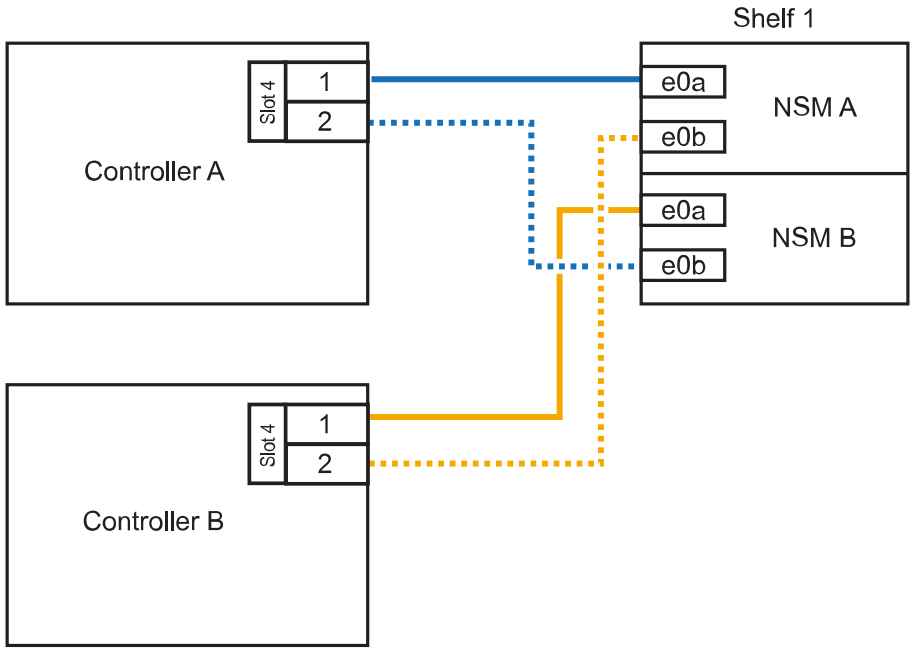

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports on each controller, complete the following substeps.

Shelves	Cabling
Shelf 1	a. Cable NSM A port e0a to controller A slot 4 port 1 (e4a). b. Cable NSM A port e0b to controller B slot 5 port 2 (e5b). c. Cable NSM B port e0a to controller B port slot 4 port 1 (e4a). d. Cable NSM B port e0b to controller A slot 5 port 2 (e5b). e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3.
Shelf 2	a. Cable NSM A port e0a to controller A slot 5 port 1 (e5a). b. Cable NSM A port e0b to controller B slot 4 port 2 (e4b). c. Cable NSM B port e0a to controller B slot 5 port 1 (e5a). d. Cable NSM B port e0b to controller A slot 4 port 2 (e4b). e. Go to step 3.

The following illustration shows cabling for two hot-added shelves:

AFF C400 HA pair with two NS224 shelves

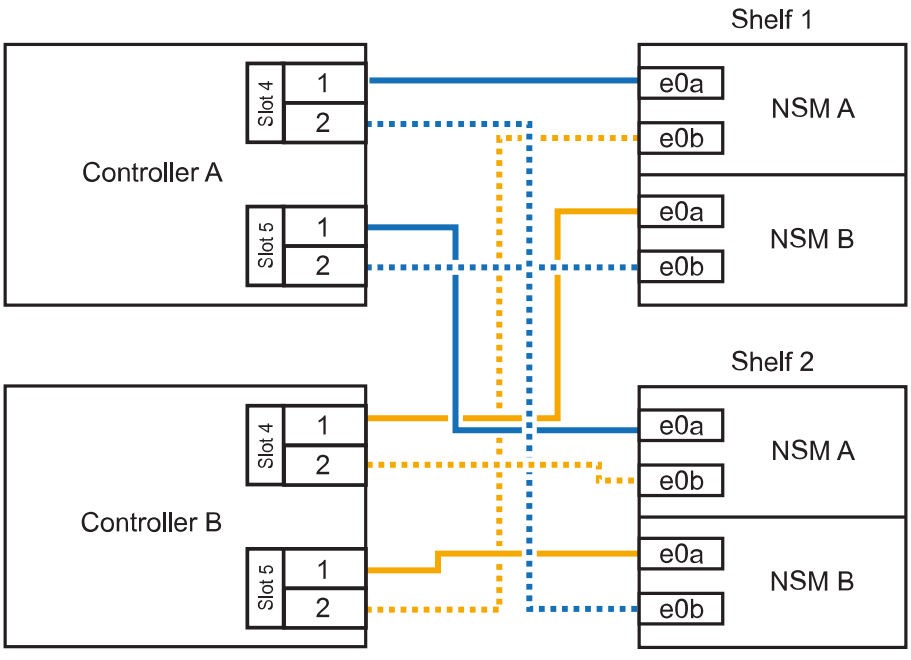

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What’s next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to ASA A800 or ASA C800

How you cable an NS224 shelf in an ASA A800 or ASA C800 HA pair depends on the number of shelves you are hot-adding and the number of RoCE-capable port sets (one or two) you are using on the controllers.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (one RoCE-capable PCIe card) on each controller, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   This step assumes you installed the RoCE-capable PCIe card in slot 5.

   a. Cable shelf NSM A port e0a to controller A slot 5 port a (e5a).
   b. Cable shelf NSM A port e0b to controller B slot 5 port b (e5b).
   c. Cable shelf NSM B port e0a to controller B slot 5 port a (e5a).
   d. Cable shelf NSM B port e0b to controller A slot 5 port b (e5b).

   The following illustration shows cabling for one hot-added shelf using one RoCE-capable PCIe card on each controller:

   AFF A800 or AFF C800 HA pair with one NS224 shelf

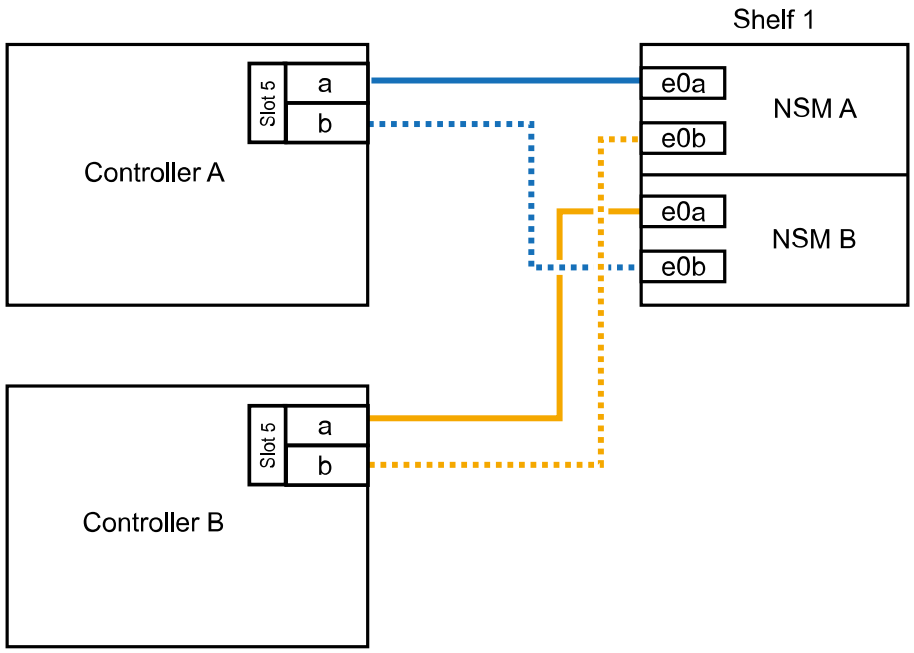

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports (two RoCE-capable

PCIe cards) on each controller, complete the applicable substeps.

This step assumes you installed the RoCE-capable PCIe cards in slot 5 and slot 3.

Shelves	Cabling
Shelf 1	These substeps assume you are beginning the cabling by cabling shelf port e0a to the RoCE-capable PCIe card in slot 5, instead of slot 3. a. Cable NSM A port e0a to controller A slot 5 port a (e5a). b. Cable NSM A port e0b to controller B slot 3 port b (e3b). c. Cable NSM B port e0a to controller B slot 5 port a (e5a). d. Cable NSM B port e0b to controller A slot 3 port b (e3b). e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3.
Shelf 2	These substeps assume you are beginning the cabling by cabling shelf port e0a to the RoCE-capable PCIe card in slot 3, instead of slot 5 (which correlates with the cabling substeps for shelf 1). a. Cable NSM A port e0a to controller A slot 3 port a (e3a). b. Cable NSM A port e0b to controller B slot 5 port b (e5b). c. Cable NSM B port e0a to controller B slot 3 port a (e3a). d. Cable NSM B port e0b to controller A slot 5 port b (e5b). e. Go to step 3.

The following illustration shows cabling for two hot-added shelves:

AFF A800 or AFF C800 HA pair with two NS224 shelves

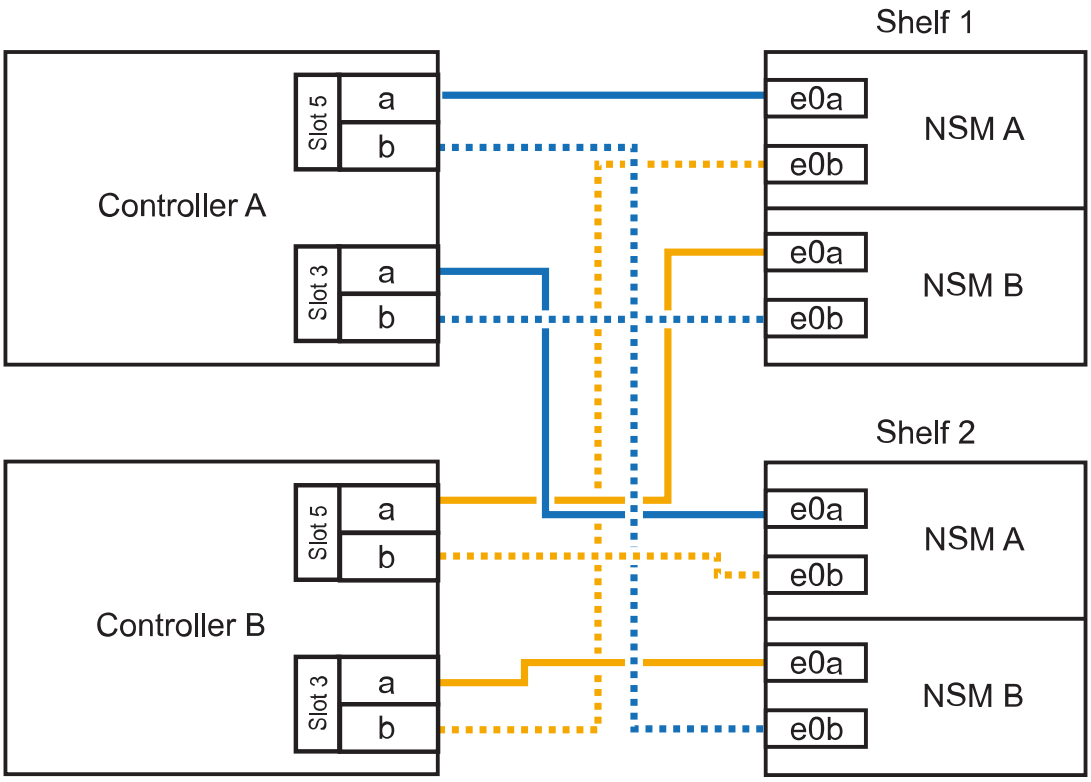

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to ASA A900

When additional storage is needed, you can hot-add up to three additional NS224 drive shelves (for a total of four shelves) to an ASA A900 HA pair.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

- This procedure assumes that your HA pair has at least one existing NS224 shelf and that you are hot-adding up to three additional shelves.
- If your HA pair has only one existing NS224 shelf, this procedure assumes that the shelf is cabled across two RoCE-capable 100GbE I/O modules on each controller.

Steps

1. If the NS224 shelf you are hot-adding will be the second NS224 shelf in the HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 10 port a (e10a).
   b. Cable shelf NSM A port e0b to controller B slot 2 port b (e2b).
   c. Cable shelf NSM B port e0a to controller B slot 10 port a (e10a).
   d. Cable shelf NSM B port e0b to controller A slot 2 port b (e2b).

      The following illustration shows the second shelf cabling (and the first shelf).

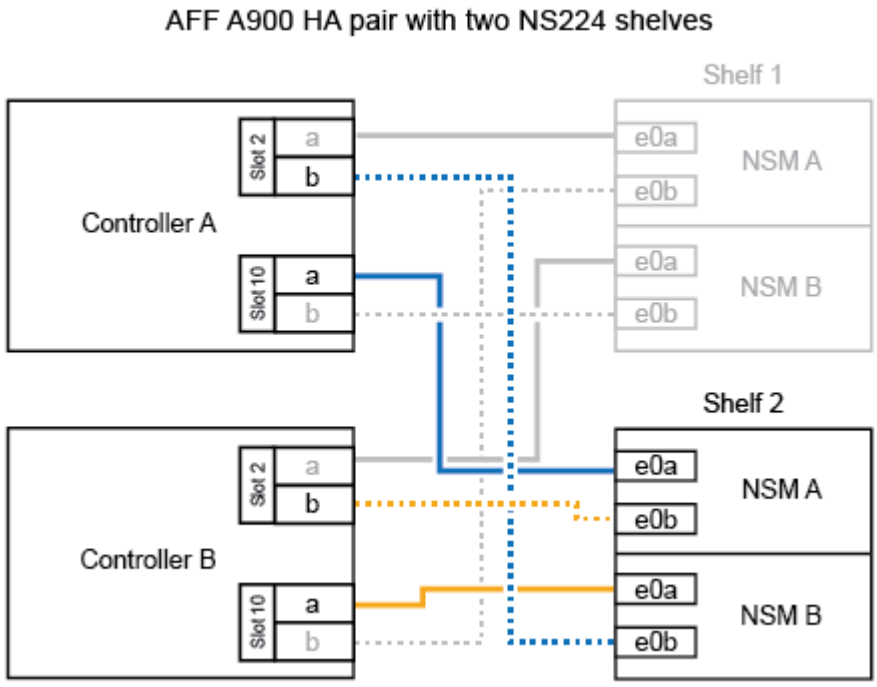

2. If the NS224 shelf you are hot-adding will be the third NS224 shelf in the HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 1 port a (e1a).
   b. Cable shelf NSM A port e0b to controller B slot 11 port b (e11b).
   c. Cable shelf NSM B port e0a to controller B slot 1 port a (e1a).
   d. Cable shelf NSM B port e0b to controller A slot 11 port b (e11b).

      The following illustration shows the third shelf cabling.

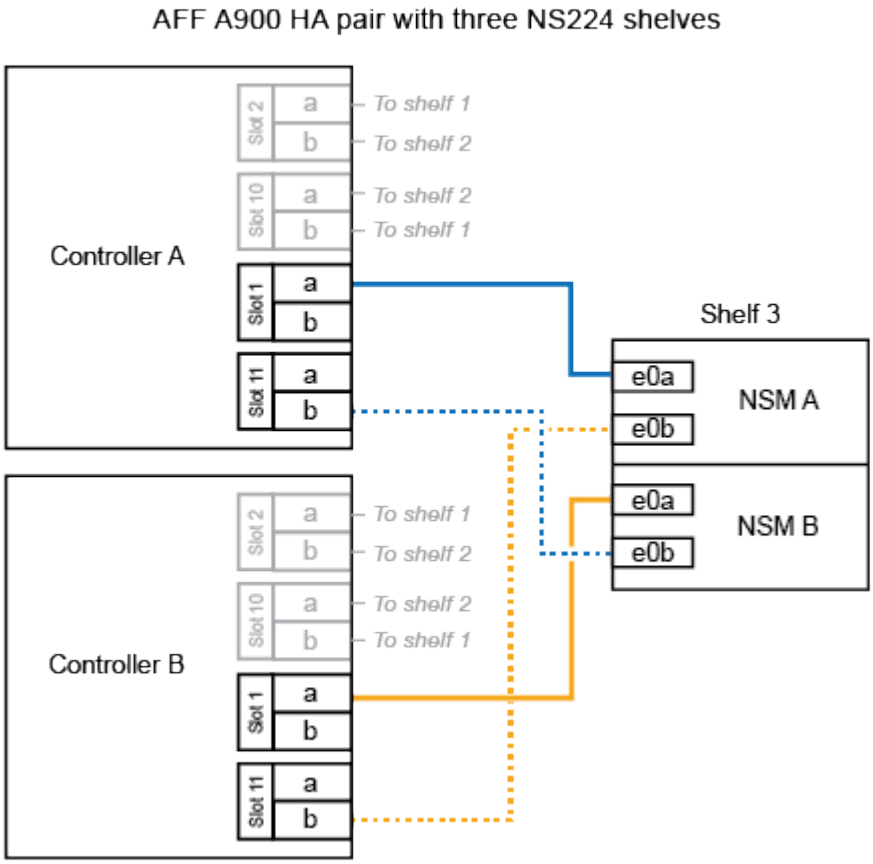

3. If the NS224 shelf you are hot-adding will be the fourth NS224 shelf in the HA pair, complete the following substeps.

   Otherwise, go to the next step.

   a. Cable shelf NSM A port e0a to controller A slot 11 port a (e11a).
   b. Cable shelf NSM A port e0b to controller B slot 1 port b (e1b).
   c. Cable shelf NSM B port e0a to controller B slot 11 port a (e11a).
   d. Cable shelf NSM B port e0b to controller A slot 1 port b (e1b).

      The following illustration shows the fourth shelf cabling.

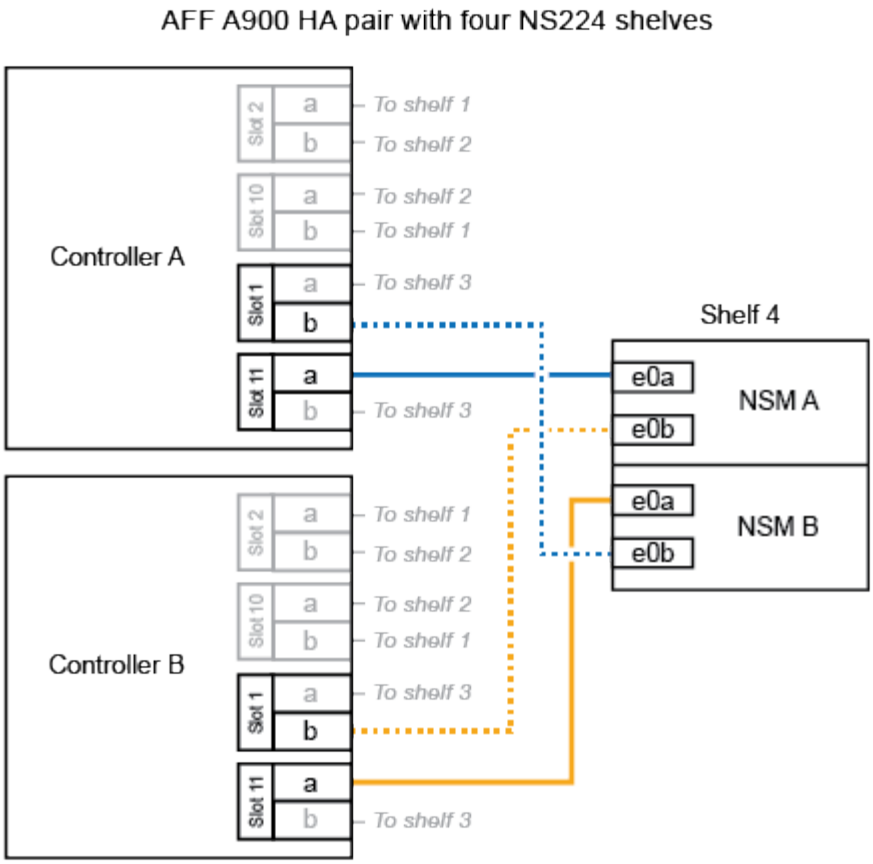

4. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to end-of-availability systems - NS224 shelves

You cable each NS224 shelf you are hot-adding so that each shelf has two connections to each controller in the HA pair.

Cable shelf to AFF A320

You can hot-add a second shelf to an existing HA pair when additional storage is needed.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

Steps

1. Cable the shelf to the controllers.
   a. Cable NSM A port e0a to controller A port e0e.
   b. Cable NSM A port e0b to controller B port e0b.
   c. Cable NSM B port e0a to controller B port e0e.
   d. Cable NSM B port e0b to controller A port e0b.

   The following illustration shows cabling for the hot-added shelf (shelf 2):

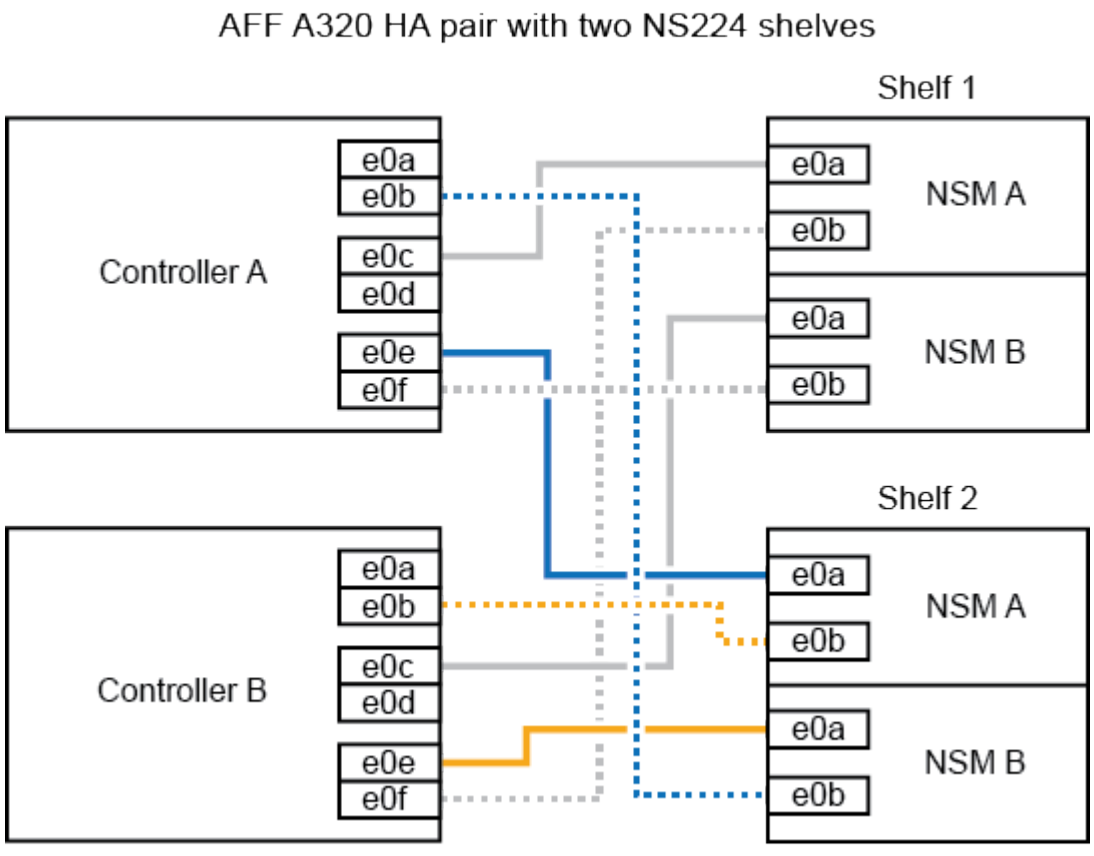

2. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to AFF A700

How you cable an NS224 shelf to an AFF A700 HA pair depends on the number of shelves you are hot-adding and the number of RoCE-capable port sets (one or two) you are using on the controllers.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.
- If you are hot-adding the initial NS224 shelf (no NS224 shelf exists in your HA pair), you must install a core dump module (X9170A, NVMe 1TB SSD) in each controller to support core dumps (store core files).

  See Replace the caching module or add/replace a core dump module — AFF A700 and FAS9000.

Steps

1. If you are hot-adding one shelf using one set of RoCE-capable ports (one RoCE capable I/O module) on each controller, and this is the only NS224 shelf in your HA pair, complete the following substeps.

   Otherwise, go to the next step.

   > This step assumes that you installed the RoCE-capable I/O module in slot 3, instead of slot 7, on each controller.

   a. Cable shelf NSM A port e0a to controller A slot 3 port a.
   b. Cable shelf NSM A port e0b to controller B slot 3 port b.
   c. Cable shelf NSM B port e0a to controller B slot 3 port a.
   d. Cable shelf NSM B port e0b to controller A slot 3 port b.

      The following illustration shows cabling for one hot-added shelf using one RoCE-capable I/O module in each controller:

AFF A700 HA pair with one NS224 shelf

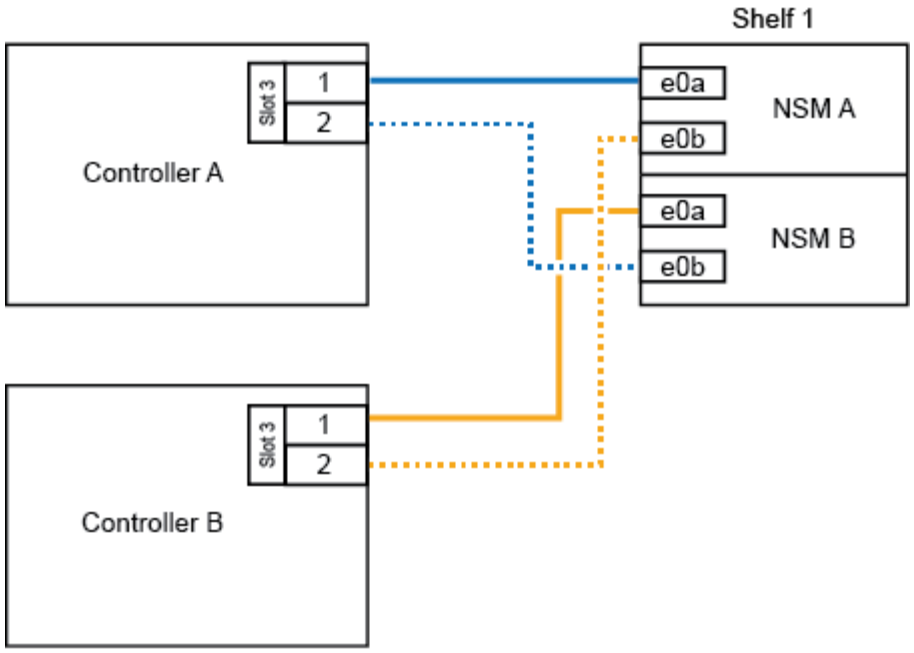

2. If you are hot-adding one or two shelves using two sets of RoCE-capable ports (two RoCE-capable I/O modules) in each controller, complete the applicable substeps.

Shelves	Cabling
Shelf 1	These substeps assume that you are beginning the cabling by cabling shelf port e0a to the RoCE-capable I/O module in slot 3, instead of slot 7. a. Cable NSM A port e0a to controller A slot 3 port a. b. Cable NSM A port e0b to controller B slot 7 port b. c. Cable NSM B port e0a to controller B slot 3 port a. d. Cable NSM B port e0b to controller A slot 7 port b. e. If you are hot-adding a second shelf, complete the “Shelf 2” substeps; otherwise, go to step 3.
Shelf 2	These substeps assume that you are beginning the cabling by cabling shelf port e0a to the RoCE-capable I/O module in slot 7, instead of slot 3 (which correlates with the cabling substeps for shelf 1). a. Cable NSM A port e0a to controller A slot 7 port a. b. Cable NSM A port e0b to controller B slot 3 port b. c. Cable NSM B port e0a to controller B slot 7 port a. d. Cable NSM B port e0b to controller A slot 3 port b. e. Go to step 3.

The following illustration shows cabling for the first and second hot-added shelves:

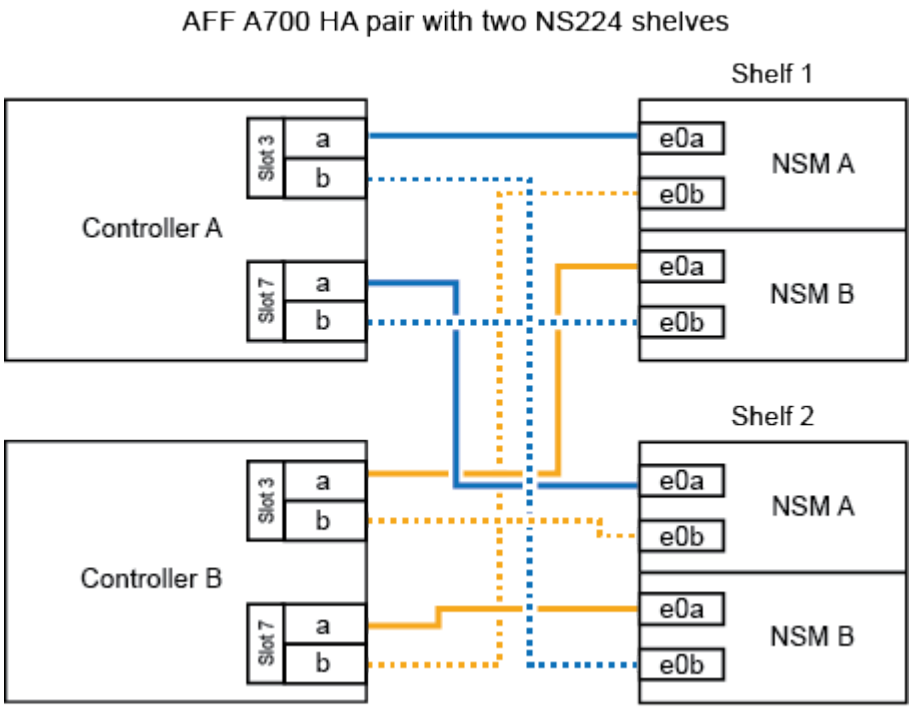

3. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Cable shelf to FAS500f

When additional storage is needed, you can hot-add one NS224 shelf to a FAS500f HA pair.

Before you begin

- You must have reviewed the hot-add requirements and best practices.
- You must have completed the applicable procedures in Prepare-to hot-add a shelf.
- You must have installed the shelves, powered them on, and set the shelf IDs as described in Install a shelf for a hot-add.

About this task

When viewed from the rear of the platform chassis, the RoCE-capable card port on the left is port "a" (e1a) and the port on the right is port "b" (e1b).

Steps

1. Cable the shelf connections:
   a. Cable shelf NSM A port e0a to controller A slot 1 port a (e1a).
   b. Cable shelf NSM A port e0b to controller B slot 1 port b (e1b).
   c. Cable shelf NSM B port e0a to controller B slot 1 port a (e1a).
   d. Cable shelf NSM B port e0b to controller A slot 1 port b (e1b).

   The following illustration shows the shelf cabling when completed.

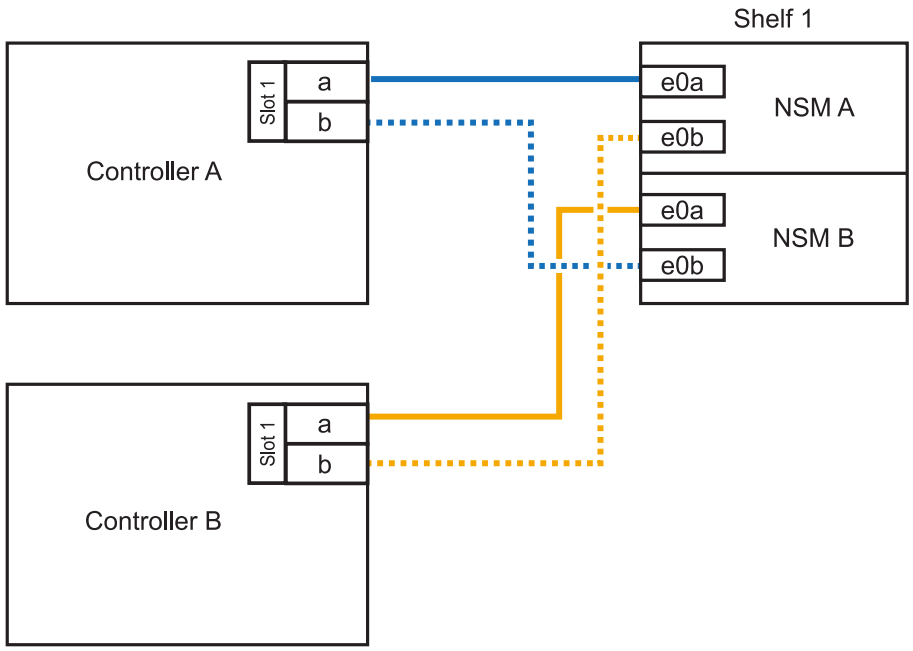

2. Verify that the hot-added shelf is cabled correctly using Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

What's next?

If you disabled automatic drive assignment as part of the preparation for this procedure, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed. Go to Complete the hot-add.

Otherwise, you are done with the hot-add shelf procedure.

Complete the hot-add - NS224 shelves

If you disabled automatic drive assignment as part of the preparation for the NS224 shelf hot-add, you need to manually assign drive ownership and then reenable automatic drive assignment, if needed.

Before you begin

You must have already cabled your shelf as instructed for your HA pair. See Overview of cabling for a hot-add.

Steps

1. Display all unowned drives: `storage disk show -container-type unassigned`

   You can enter the command on either controller.

2. Assign each drive: `storage disk assign -disk *disk_name* -owner *owner_name*`

   You can enter the command on either controller.

   You can use the wild card character to assign more than one drive at once.

3. Reenable automatic drive assignment if needed: `storage disk option modify -node *node_name* -autoassign on`

   You must reenable automatic drive assignment on both controllers.

Change a shelf ID - NS224 shelves

You can change a shelf ID in a system when ONTAP is not yet running or when hot-adding a shelf prior to it being cabled to the system. You can also change a shelf ID when ONTAP is up and running (controller modules are available to serve data) and all drives in the shelf are unowned, spares, or part of offlined aggregate(s).

Before you begin

- If ONTAP is up and running (controller modules are available to serve data), you must have verified that all drives in the shelf are unowned, spares, or part of offlined aggregate(s).

  You can verify the state of the drives by using the `storage disk show -shelf *shelf_number*` command. Output in the `Container Type` column should display `spare` or `broken` if it is a failed drive. Additionally, the `Container Name` and `Owner` columns should have a dash.

- You need a paper clip with one side straightened or a narrow-tipped ballpoint pen.

  You use the paper clip or ballpoint pen to access the shelf ID button through the small hole, to the right of the LEDs, in the Operator Display Panel (ODP).

About this task

- A valid shelf ID is 00 through 99.

- Shelf IDs must be unique within an HA pair.
- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, and then plug them back in) in order for the shelf ID to take effect.

  The amount of time you wait before plugging the power cords back in depends on the state of ONTAP, as described later in this procedure.

  NS224 shelves do not have power switches on the power supplies.

Steps

1. Power on the shelf, if it's not already on.

   You connect the power cords first to the shelf, securing them in place with the power cord retainer, and then connect the power cords to different power sources for resiliency.

2. Remove the left end cap to locate the small hole to the right of the LEDs.

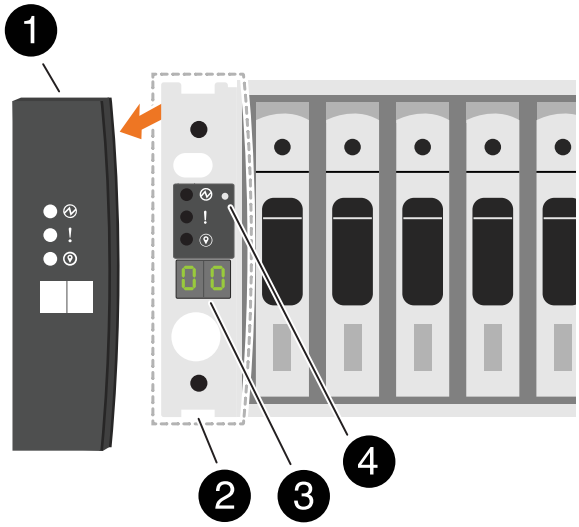

1	Shelf end cap
2	Shelf faceplate
3	Shelf ID number
4	Shelf ID button

3. Change the first number of the shelf ID:
   a. Insert the paper clip or ballpoint pen into the small hole.
   b. Press and hold the button until the first number on the digital display blinks, and then release the button.

      It can take up to 15 seconds for the number to blink. This activates the shelf ID programming mode.

If the ID takes longer than 15 seconds to blink, press and hold the button again, making sure to press it in all the way.

c. Press and release the button to advance the number until you reach the desired number from 0 to 9.

   Each press and release duration can be as short as one second.

   The first number continues to blink.

4. Change the second number of the shelf ID:

   a. Press and hold the button until the second number on the digital display blinks.

      It can take up to three seconds for the number to blink.

      The first number on the digital display stops blinking.

   b. Press and release the button to advance the number until you reach the desired number from 0 to 9.

      The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the button until the second number stops blinking.

   It can take up to three seconds for the number to stop blinking.

   Both numbers on the digital display start blinking and the amber LED on the ODP illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf to make the shelf ID take effect.

   You must unplug the power cord from both power supplies on the shelf, wait the appropriate amount of time, and then plug them back into the shelf power supplies to complete the power cycle.

   A power supply is powered on as soon as the power cord is plugged in. Its bicolored LED should illuminate green.

   - If ONTAP is not yet running or you are hot-adding a shelf (that has not yet been cabled to the system), wait at least 10 seconds.
   - If ONTAP is running (controllers are available to serve data), and all drives in the shelf are unowned, spares, or part of offlined aggregate(s), wait at least 70 seconds.

     This time allows ONTAP to properly delete the old shelf address and update the copy of the new shelf address.

7. Replace the left end cap.

Cable shelves as switch-attached storage - NS224 shelves

If you have a system in which the NS224 drive shelves need to be cabled as switch-attached storage (not direct-attached storage), use the information provided.

- To cable NS224 drive shelves through storage switches, you can refer to the NetApp Hardware Universe for more information. For older hardware models, switch cabling information can also be found in the

NS224 NVMe Drive Shelf Cabling Guide.

- To install your storage switches, refer to the AFF and FAS Switch Documentation.
- To confirm supported hardware for your platform model, such as storage switches and cables, refer to the NetApp Hardware Universe.

Maintain

Replace the boot media - NS224 shelves

You can replace a failed boot media in an NS224 shelf. Replacing the boot media can be done nondisruptively, while the drive shelf is powered on, and I/O is in progress.

Before you begin

- **NS224 with NSM100 modules only**: Your HA pair must already be running ONTAP 9.7 or later, which has the minimum supported version of NSM firmware.

  You can enter the `storage shelf show -module` command at the console of either controller to verify the version of NSM firmware on your shelf.

  > If your shelf is not running NSM firmware version 0111 or later, you cannot replace the boot media, you must replace the NSM.
  >
  > Replace an NSM - NS224 shelves

- **NS224 with NSM100 modules only**: You need a Phillips #1 screwdriver.

  The screw used to secure the boot media to the board requires a Phillips #1 screwdriver; using a different type of screwdriver could strip the screw.

- The shelf's partner NSM must be up and running, and be cabled correctly so that your shelf maintains connectivity when you remove the NSM with the failed FRU (target NSM).

  NetApp Downloads: Config Advisor

- All other components in the system must be functioning properly.

About this task

- After the boot media is replaced, the boot image from the shelf's partner NSM is automatically copied to the replacement boot media.

  This can take up to five minutes.

- Allow at least 70 seconds between removal and installation of the NVMe shelf module (NSM).

  This allows enough time for ONTAP to process the NSM removal event.

- If needed, you can turn on the shelf's location (blue) LEDs to aid in physically locating the affected shelf: `storage shelf location-led modify -shelf-name` *`shelf_name`* `-led-status on`

  If you do not know the *`shelf_name`* of the affected shelf, run the `storage shelf show` command.

  A shelf has three location LEDs: one on the operator display panel and one on each NSM. Location LEDs

remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the `off` option.

- After replacing the boot media, you can return the failed part to NetApp as described in the RMA instructions shipped with the kit.

  If you need the RMA number or additional help with the replacement procedure, contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific).

NSM100 modules

You can use the following animation or the written steps to replace the boot media.

Replace the NS224 shelf boot media

Steps

1. Make sure that both NSMs in the shelf are running the same version of firmware: version 0200 or later.
2. Properly ground yourself.
3. Disconnect the cabling from the NSM that contains the FRU that you are replacing:
   a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling from the NSM ports.

      Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports when you reinsert the NSM, later in this procedure.

4. Remove the NSM from the shelf:
   a. Loop your index fingers through the finger holes of the latching mechanisms on either side of the NSM.

      If you are removing the bottom NSM, and if the bottom rail is obstructing access to the latching mechanisms, place your index fingers through the finger holes from the inside (by crossing your arms).

   b. With your thumbs, press down and hold the orange tabs on top of the latching mechanisms.

      The latching mechanisms raise, clearing the latching pins on the shelf.

   c. Gently pull until the NSM is about one third of the way out of the shelf, grasp the NSM sides with both hands to support its weight, and then place it on a flat stable surface.

      When you begin pulling, the latching mechanism arms extend from the NSM and lock in their fully extended position.

5. Loosen the NSM cover thumb screw and open the cover.
6. Physically locate the failed boot media.

   The boot media is located along the shelf chassis wall opposite from the power supply.

7. Replace the boot media:
   a. Using the Phillips #1 screwdriver, carefully remove the screw securing the bottom (notched) end of the boot media to the board.
   b. Remove the boot media by rotating the notched end up slightly and then gently pulling it towards you until it releases from the socket.

You can hold the boot media by placing your thumb and forefinger on the side edges, at the notched end.

c. Unpack the boot media from the antistatic bag.

d. Insert the replacement boot media by pushing it gently into the socket until it is seated squarely and completely in the socket.

You can hold the boot media by placing your thumb and forefinger on the side edges, at the notched end. Make sure that the side with the heat sink is facing up.

When correctly seated, and when you let go of the boot media, the notched end of the boot media is angled up, away from the board, because it is not yet secured with the screw.

e. Gently hold down the notched end of the boot media as you insert and tighten the screw with the screwdriver to secure the boot media in place.

Tighten the screw just enough to hold the boot media securely in place, but do not overtighten.

8. Close the NSM cover, and then tighten the thumb screw.
9. Reinsert the NSM into the shelf:
   a. Make sure that the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, gently slide the NSM into the shelf until the weight of the NSM is fully supported by the shelf.
   c. Push the NSM into the shelf until it stops (about half an inch from the back of the shelf).

   You can place your thumbs on the orange tabs on the front of each finger loop (of the latching mechanism arms) to push in the NSM.

   d. Loop your index fingers through the finger holes of the latching mechanisms on either side of the NSM.

   If you are inserting the bottom NSM, and if the bottom rail is obstructing access to the latching mechanisms, place your index fingers through the finger holes from the inside (by crossing your arms).

   e. With your thumbs, press down and hold the orange tabs on top of the latching mechanisms.
   f. Gently push forward to get the latches over the stop.
   g. Release your thumbs from the tops of the latching mechanisms, and then continue pushing until the latching mechanisms snap into place.

   The NSM should be fully inserted into the shelf and flush with the edges of the shelf.

10. Reconnect the cabling to the NSM:
    a. Reconnect the storage cabling to the same two NSM ports.

    Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

    b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power

supply.

When functioning correctly, a power supply's bicolored LED illuminates green.

Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

11. Verify that the attention (amber) LEDs on the NSM containing the failed boot media and the shelf operator display panel are no longer illuminated.

    It can take between 5 to 10 minutes for the attention LEDs to turn off. This is the amount of time it takes the NSM to reboot and the boot media image copy to complete.

    If the fault LEDs remain on, the boot media might not be seated correctly or there might be another issue and you should contact technical support for assistance.

12. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

    If any cabling errors are generated, follow the corrective actions provided.

    NetApp Downloads: Config Advisor

NSM100B modules

You can use the following steps to replace the failed boot media.

Steps

1. Properly ground yourself.
2. Disconnect the cabling from the NSM that contains the FRU that you are replacing:
   a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling from the NSM ports.

      Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports when you reinsert the NSM, later in this procedure.
3. Remove the NSM:

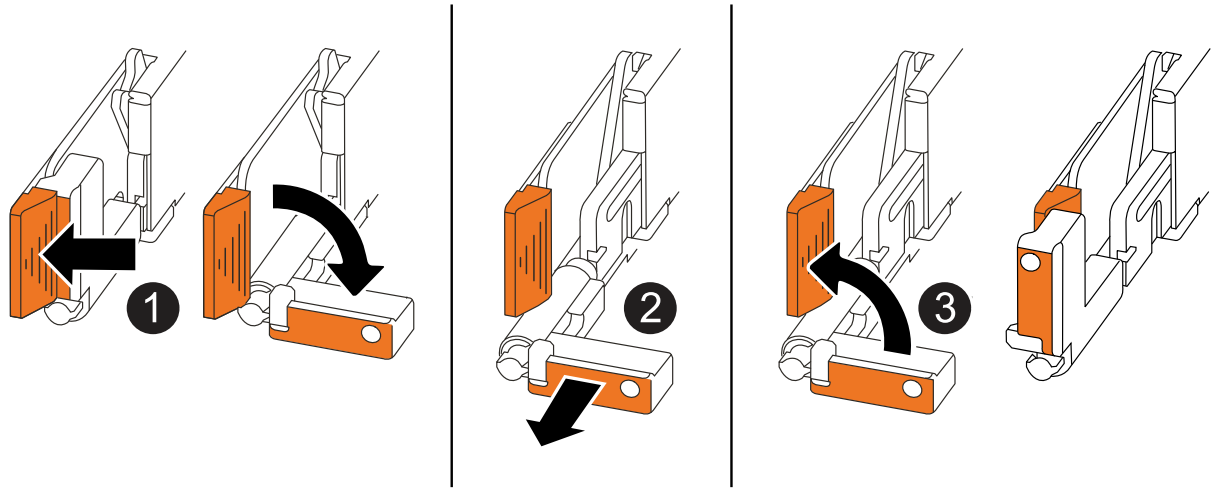

1	On both ends of the NSM, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the NSM from the midplane. As you pull, the handles extend out from the shelf. When you feel some resistance, keep pulling. • Slide the NSM out of the shelf and place it on a flat, stable surface. Make sure that you support the bottom of the NSM as you slide it out of the shelf.
3	Rotate the handles upright (next to the tabs) to move them out of the way.

4. Open the NSM cover by turning the thumbscrew counterclockwise to loosen it, and then open the cover.
5. Physically locate the failed boot media.
6. Remove the boot media:

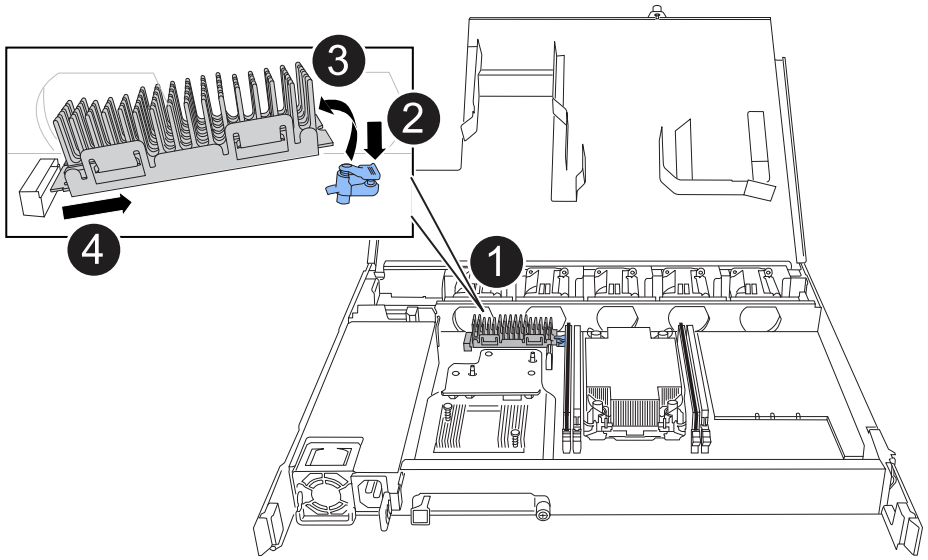

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

7. Install the replacement boot media:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the locking button.
   c. Push the locking button, rotate the boot media all the way down, and then release the locking button.
8. Close the NSM cover, and then tighten the thumb screw.
9. Insert the NSM into the shelf:

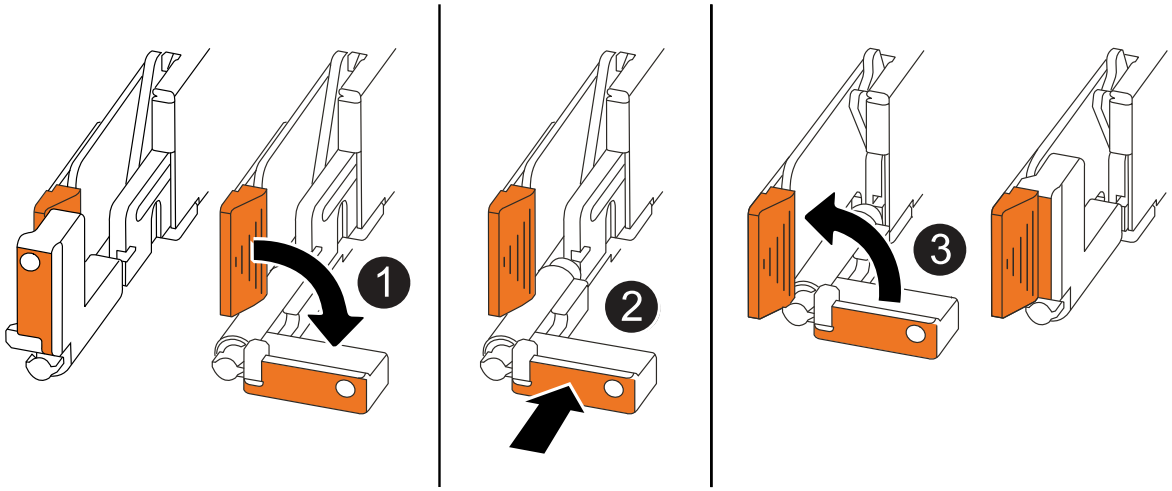

1	If you rotated the NSM handles upright (next to the tabs) to move them out of the way while you serviced the NSM, rotate them down to the horizontal position.
2	Align the rear of the NSM with the opening in the shelf, and then gently push the NSM using the handles until it is fully seated.
3	Rotate the handles to the upright position and lock in place with the tabs.

10. Reconnect the cabling to the NSM:
    a. Reconnect the storage cabling to the same two NSM ports.

       Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

    b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power supply.

       When functioning correctly, a power supply's bicolored LED illuminates green.

       Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

11. Verify that the attention (amber) LEDs on the NSM containing the failed boot media and the shelf operator display panel are no longer illuminated.

    It can take between 5 to 10 minutes for the attention LEDs to turn off. This is the amount of time it takes the NSM to reboot and the boot media image copy to complete.

    If the fault LEDs remain on, the boot media might not be seated correctly or there might be another issue and you should contact technical support for assistance.

12. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

    If any cabling errors are generated, follow the corrective actions provided.

NetApp Downloads: Config Advisor

Replace a DIMM - NS224 shelves

You can replace a faulty DIMM nondisruptively in an NS224 drive shelf that is powered on, and while I/O is in progress.

Before you begin

- The shelf's partner NSM must be up and running, and be cabled correctly so that your shelf maintains connectivity when you remove the NSM with the failed FRU (target NSM).

  NetApp Downloads: Config Advisor

- All other components in the system, including the other three DIMMs in the NSM100 module and one DIMM in the NSM100B module, must be functioning properly.

About this task

- Allow at least 70 seconds between removal and installation of the NVMe shelf module (NSM).

  This allows enough time for ONTAP to process NSM removal event.

- **Best practice:** The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.

  NetApp Downloads: Disk Shelf Firmware

  NetApp Downloads: Disk Drive Firmware

  Do not revert firmware to a version that does not support your shelf and its components.

- If needed, you can turn on the shelf's location (blue) LEDs to aid in physically locating the affected shelf: `storage shelf location-led modify -shelf-name shelf_name -led-status on`

  If you do not know the *`shelf_name`* of the affected shelf, run the `storage shelf show` command.

  A shelf has three location LEDs: one on the operator display panel and one on each NSM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the `off` option.

- When you unpack the replacement DIMM, save all packing materials for use when you return the failed DIMM.

  If you need the RMA number or additional help with the replacement procedure, contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific).

NSM100 modules

You can use the following animation or the written steps to replace a DIMM.

Replace a DIMM in an NS224 shelf

Steps

1. Properly ground yourself.
2. Disconnect the cabling from the NSM that contains the FRU that you are replacing:
   a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling from the NSM ports.

      Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports when you reinsert the NSM, later in this procedure.

3. Remove the NSM from the shelf:
   a. Loop your index fingers through the finger holes of the latching mechanisms on either side of the NSM.

      If you are removing the bottom NSM, and if the bottom rail is obstructing access to the latching mechanisms, place your index fingers through the finger holes from the inside (by crossing your arms).

   b. With your thumbs, press down and hold the orange tabs on top of the latching mechanisms.

      The latching mechanisms raise, clearing the latching pins on the shelf.

   c. Gently pull until the NSM is about one third of the way out of the shelf, grasp the NSM sides with both hands to support its weight, and then place it on a flat stable surface.

      When you begin pulling, the latching mechanism arms extend from the NSM and lock in their fully extended position.

4. Loosen the NSM cover thumb screw and open the cover.

   The FRU label on the NSM cover shows the location of the four DIMMs, two on either side of the heat sink, in the center of the NSM.

5. Physically identify the faulty DIMM.

   When a DIMM is faulty, the system logs a warning message to the system console indicating which DIMM is faulty.

6. Replace the faulty DIMM:
   a. Note the orientation of the DIMM in the slot so that you can insert the replacement DIMM using the same orientation.
   b. Eject the DIMM from its slot by slowly pushing apart the ejector tabs at both ends of the DIMM

slot, and then lift the DIMM out of the slot.

> Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.

The ejector tabs remain in the open position.

c. Remove the replacement DIMM from its antistatic shipping bag.

d. Hold the DIMM by the corners, and then insert the DIMM squarely into a slot.

The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

When inserted correctly, the DIMM should go in easily but fit tightly in the slot. If not, reinsert the DIMM.

e. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

7. Close the NSM cover, and then tighten the thumb screw.

8. Reinsert the NSM into the shelf:

a. Make sure that the latching mechanism arms are locked in the fully extended position.

b. Using both hands, gently slide the NSM into the shelf until the weight of the NSM is fully supported by the shelf.

c. Push the NSM into the shelf until it stops (about half an inch from the back of the shelf).

You can place your thumbs on the orange tabs on the front of each finger loop (of the latching mechanism arms) to push in the NSM.

d. Loop your index fingers through the finger holes of the latching mechanisms on either side of the NSM.

> If you are inserting the bottom NSM, and if the bottom rail is obstructing access to the latching mechanisms, place your index fingers through the finger holes from the inside (by crossing your arms).

e. With your thumbs, press down and hold the orange tabs on top of the latching mechanisms.

f. Gently push forward to get the latches over the stop.

g. Release your thumbs from the tops of the latching mechanisms, and then continue pushing until the latching mechanisms snap into place.

The NSM should be fully inserted into the shelf and flush with the edges of the shelf.

9. Reconnect the cabling to the NSM:

a. Reconnect the storage cabling to the same two NSM ports.

Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power supply.

When functioning correctly, a power supply's bicolored LED illuminates green.

Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

10. Verify that the attention (amber) LEDs on the NSM containing the failed DIMM and the shelf operator display panel are no longer illuminated.

    The NSM attention LEDs turn off after the NSM reboots and no longer detects a DIMM issue. This can take three to five minutes.

11. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

    If any cabling errors are generated, follow the corrective actions provided.

    NetApp Downloads: Config Advisor

NSM100B modules

Steps

1. Properly ground yourself.
2. Disconnect the cabling from the NSM that contains the FRU that you are replacing:
   a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling from the NSM ports.

      Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports when you reinsert the NSM, later in this procedure.
3. Insert the NSM into the shelf:

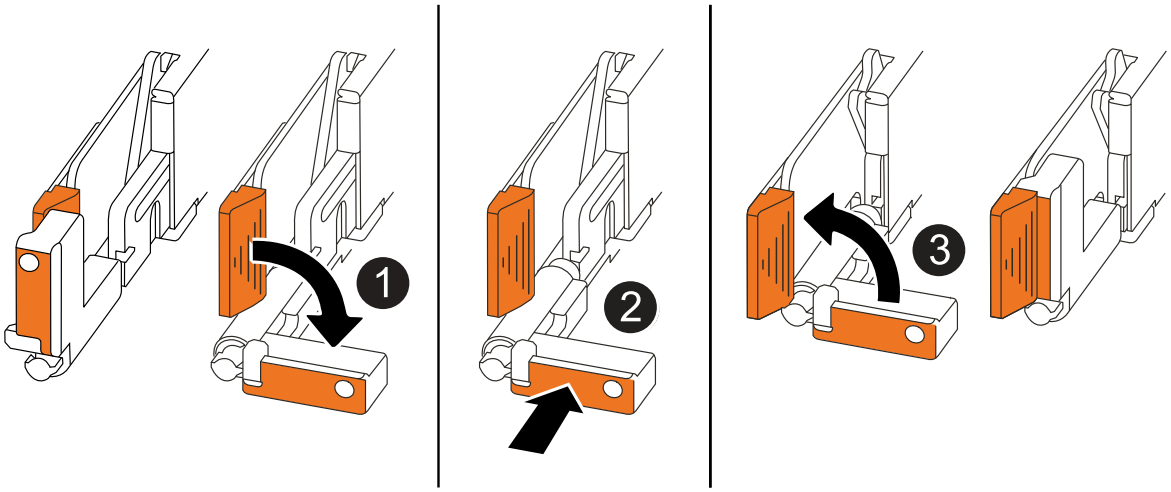

1	If you rotated the NSM handles upright (next to the tabs) to move them out of the way while you serviced the NSM, rotate them down to the horizontal position.
2	Align the rear of the NSM with the opening in the shelf, and then gently push the NSM using the handles until it is fully seated.
3	Rotate the handles to the upright position and lock in place with the tabs.

4. Open the NSM cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

   The FRU label on the NSM cover shows the location of the two DIMMs and two DIMM Blanks in the NSM.

5. Physically identify the faulty DIMM.

   When a DIMM is faulty, the system logs a warning message to the system console indicating which DIMM needs to be replaced.

6. Remove the faulty DIMM:

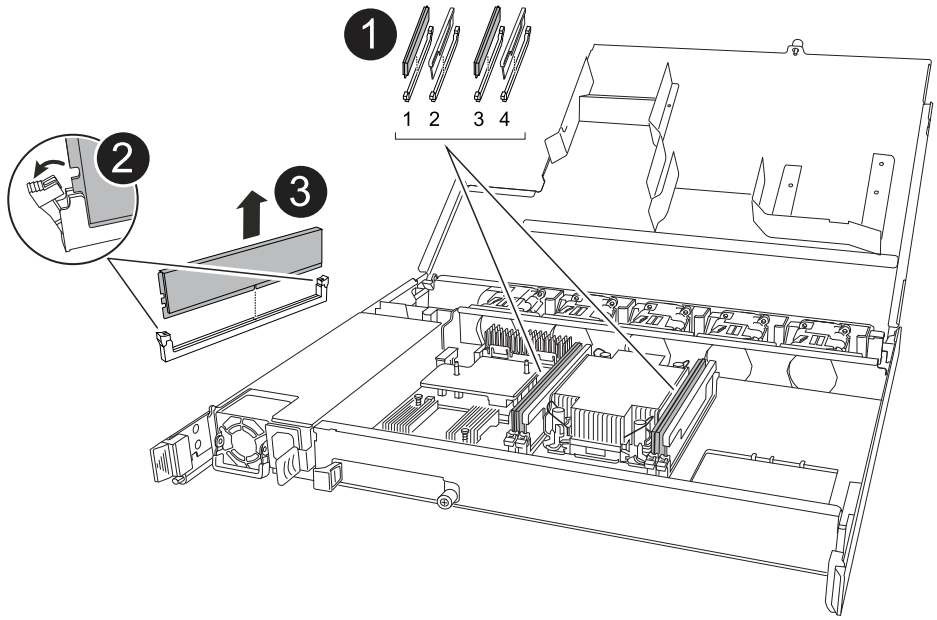

1	DIMM slot numbering and positions. The NSM contains DIMMs in slots 1 and 3, and DIMM Blanks in slots 2 and 4.

2	• Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM using the same orientation. • Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

7. Replace the DIMM:
   a. Remove the replacement DIMM from its antistatic shipping bag.
   b. Hold the DIMM by the corners, and then insert the DIMM squarely into a slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

      When inserted correctly, the DIMM should go in easily but fit tightly in the slot. If not, reinsert the DIMM.

   c. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.
8. Close the NSM cover, and then tighten the thumb screw.
9. Insert the NSM into the shelf:

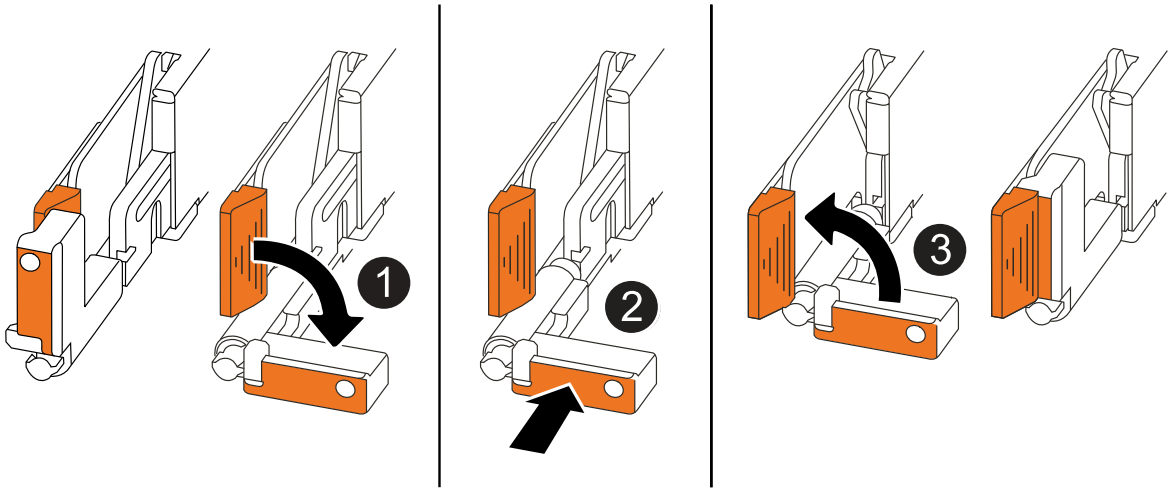

1	If you rotated the NSM handles upright (next to the tabs) to move them out of the way while you serviced the NSM, rotate them down to the horizontal position.
2	Align the rear of the NSM with the opening in the shelf, and then gently push the NSM using the handles until it is fully seated.

3	Rotate the handles to the upright position and lock in place with the tabs.

10. Reconnect the cabling to the NSM:
    a. Reconnect the storage cabling to the same two NSM ports.

       Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

    b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power supply.

       When functioning correctly, a power supply's bicolored LED illuminates green.

       Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

11. Verify that the attention (amber) LEDs on the NSM containing the failed DIMM and the shelf operator display panel are no longer illuminated.

    The NSM attention LEDs turn off after the NSM reboots and no longer detects a DIMM issue. This can take three to five minutes.

12. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

    If any cabling errors are generated, follow the corrective actions provided.

    NetApp Downloads: Config Advisor

Hot-swap a drive - NS224 shelves

You can replace a failed drive nondisruptively in an NS224 drive shelf that is powered on, and while I/O is in progress.

Before you begin

- The drive that you are installing must be supported by the NS224 shelf.

  NetApp Hardware Universe

- If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP documentation.

  Instructions in the ONTAP documentation describe additional steps you must perform before and after replacing an SED.

  NetApp encryption overview with the CLI

- All other components in the system must be functioning properly; if not, contact technical support.
- Verify that the drive you are removing is failed.

  You can verify that the drive is failed by running the `storage disk show -broken` command. The

failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

> Depending on the drive type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

About this task

- **Best practice:** The best practice is to have the current version of the Disk Qualification Package (DQP) installed before hot-swapping a drive.

  Having the current version of the DQP installed allows your system to recognize and use newly qualified drives. This avoids system event messages about having noncurrent drive information and prevention of drive partitioning because drives are not recognized. The DQP also notifies you of noncurrent drive firmware.

  NetApp Downloads: Disk Qualification Package

- **Best practice:** The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.

  NetApp Downloads: Disk Shelf Firmware

  NetApp Downloads: Disk Drive Firmware

  > Do not revert firmware to a version that does not support your shelf and its components.

- Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

  > Drive firmware checks occur every two minutes.

- If needed, you can turn on the shelf's location (blue) LEDs to aid in physically locating the affected shelf: `storage shelf location-led modify -shelf-name shelf_name -led-status on`

  If you do not know the *`shelf_name`* of the affected shelf, run the `storage shelf show` command.

  A shelf has three location LEDs: one on the operator display panel and one on each NSM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the `off` option.

- When you unpack the replacement drive, save all packing materials for use when you return the failed drive.

  If you need the RMA number or additional help with the replacement procedure, contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific).

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment if it is enabled.

You need to manually assign drive ownership if drives in the shelf are owned by both controller modules in the HA pair. You complete this task later in this procedure.

a. Verify whether automatic drive assignment is enabled: `storage disk option show`

   You can enter the command on either controller module.

   If automatic drive assignment is enabled, the output shows `on` in the `Auto Assign` column (for each controller module).

b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

   You must disable automatic drive assignment on both controller modules.

2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

4. Remove the failed drive:

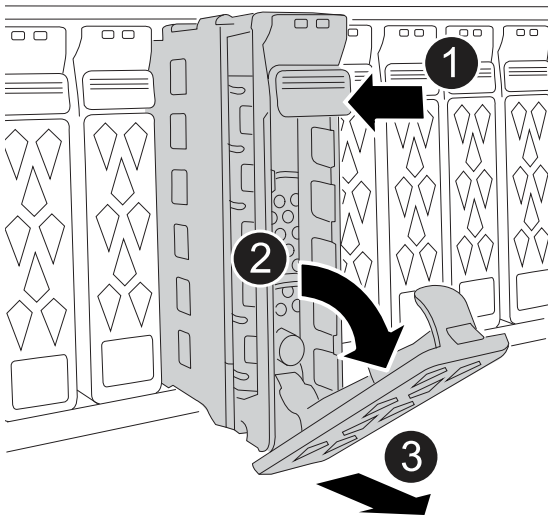

1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.

3	Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.

5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the drive.
   b. Gently push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the mid plane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

7. Verify that the drive's activity (green) LED is illuminated.

   When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat steps 3 through step 7.
9. If you disabled automatic drive assignment in step 1, manually assign drive ownership, and then reenable automatic drive assignment if needed:
   a. Display all unowned drives: `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

      You must reenable automatic drive assignment on both controller modules.

Drive shelf

Overview of shelf maintenance - NS224 shelves

You can take the following actions to maintain your NS224 shelf:

- Hot-add a drive
- Cold-replace a shelf
- Hot-remove a shelf
- Monitor shelf LEDs

Cold-replace a shelf - NS224 shelves

When replacing a drive shelf in a production system that has disks in use, you need to perform a cold shelf replacement. This is a disruptive procedure. It requires you to halt the controllers in your HA pair.

Use the NetApp Knowledge Base article How to replace a shelf chassis using a cold shelf removal procedure.

Hot-add a drive - NS224 shelves

You can add new drives to a powered-on shelf non-disruptively, even during I/O operations.

Use the NetApp Knowledge Base article Best practices for adding disks to an existing shelf or cluster.

Hot-remove a shelf - NS224 shelves

You can hot-remove an NS224 drive shelf that has had the aggregates removed from the drives, in an HA pair that is up and serving data (I/O is in progress).

This procedure does not apply to ASA r2 systems.

Before you begin

- Your HA pair cannot be in a takeover state.
- You must have removed all aggregates from the drives (the drives must be spares) in the shelf you are removing.

  If you attempt this procedure with aggregates on the shelf you are removing, you could fail the system with a multidisk panic.

  You can use the `storage aggregate offline -aggregate` *`aggregate_name`* command and then the `storage aggregate delete -aggregate` *`aggregate_name`* command.

  To review more information on this step and avoid potential IO issues, see the Disks and aggregates overview.

- If your system shipped in a system cabinet, you need a Phillips screwdriver to remove the screws securing the shelf to the cabinet rack rails.

About this task

- If you are hot-removing more than one shelf, you remove one shelf at a time.
- **Best practice:** The best practice is to clear drive ownership after you remove the aggregates from the drives in the shelf you are removing.

  Clearing ownership information from a spare drive allows the drive to be properly integrated into another node (as needed).

  The procedure for removing ownership from drives can be found in the disks and aggregates content:

  Disks and aggregates overview

The procedure requires you to disable automatic drive assignment. You reenable automatic drive assignment at the end of this procedure (after you have hot-removed the shelf).

- If needed, you can turn on the shelf's location (blue) LEDs to aid in physically locating the affected shelf: `storage shelf location-led modify -shelf-name shelf_name -led-status on`

  If you do not know the *shelf_name* of the affected shelf, run the `storage shelf show` command.

  A shelf has three location LEDs: one on the operator display panel and one on each NSM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the off option.

- After disconnecting a shelf from non-dedicated RoCE capable ports (on board the controllers, on RoCE capable PCIe cards, a combination of both, or on I/O modules), you have the option of reconfiguring these ports for networking use.

  If your HA pair is running ONTAP 9.7 or later, you do not need to reboot the controllers, unless one or both controllers are in maintenance mode. This procedure assumes that neither controller is in maintenance mode.

Steps

1. Properly ground yourself.
2. Verify that the drives in the shelf you are removing have no aggregates (are spares) and that ownership is removed:

   a. Enter the following command to list all of the drives in the shelf that you are removing: `storage disk show -shelf shelf_number`

      You can enter the command on either controller module.

   b. Check the output to verify that there are no aggregates on the drives.

      Drives with no aggregates have a dash in the `Container Name` column.

   c. Check the output to verify that ownership is removed from the drives.

      Drives with no ownership have a dash in the `Owner` column.

      If you have failed drives, they display broken in the `Container Type` column. (Failed drives do not have ownership.)

      The following output shows drives on the shelf being removed (shelf 2) are in a correct state for removing the shelf. The aggregates are removed on all of the drives; therefore, a dash appears in the `Container Name` column for each drive. Ownership is also removed on all of the drives; therefore, a dash appears in the `Owner` column for each drive.

```
cluster1::> storage disk show -shelf 2

           Usable           Disk     Container   Container
Disk         Size Shelf Bay Type     Type        Name       Owner
-------- -------- ----- --- ------   ----------- ---------- ---------
...
2.2.4           -     2   4 SSD-NVM  spare                  -         -
2.2.5           -     2   5 SSD-NVM  spare                  -         -
2.2.6           -     2   6 SSD-NVM  broken                 -         -
2.2.7           -     2   7 SSD-NVM  spare                  -         -
...
```

3. Physically locate the shelf you are removing.
4. Disconnect the cabling from the shelf you are removing:
   a. Disconnect the power cords from the power supplies by opening the power cord retainer if they are AC power supplies, or unscrewing the two thumb screws if they are DC power supplies, and then unplug the power cords from the power supplies.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling (from the shelf to the controllers).
5. Physically remove the shelf from the rack or cabinet.

   A fully loaded NS224 shelf can weigh up to 66.78 lbs (30.29 kg) with NSM100 modules or an average of 56.8 lbs (25.8 kg) with NSM100B modules and requires two people to lift or use of a hydraulic lift. Avoid removing shelf components (from the front or rear of the shelf) to reduce the shelf weight, because shelf weight will become unbalanced.

   If your system was shipped in a cabinet, you must first unscrew the two Phillips screws securing the shelf to the rack rails. The screws are located on the inside shelf walls of the bottom NSM. You should remove both NSMs to access the screws.

6. If you are removing more than one shelf, repeat steps 2 through 5.

   Otherwise, go to the next step.

7. If you disabled automatic drive assignment when you removed ownership from the drives, reenable it: `storage disk option modify -autoassign on`

   You run the command on both controller modules.

8. You have the option of reconfiguring the non-dedicated RoCE capable ports for networking use, by completing the following substeps. Otherwise, you are done with this procedure.
   a. Verify the names of the non-dedicated ports, currently configured for storage use: `storage port show`

      You can enter the command on either controller module.

The non-dedicated ports configured for storage use are displayed in the output as follows:
If your HA pair is running ONTAP 9.8 or later, the non-dedicated ports display `storage` in the `Mode` column.
If your HA pair is running ONTAP 9.7, the non-dedicated ports, which display `false` in the `Is Dedicated?` column, also display `enabled` in the `State` column.

b. Complete the set of steps applicable to the version of ONTAP your HA pair is running:

If your HA pair is running…	Then…
ONTAP 9.8 or later	a. Reconfigure the non-dedicated ports for networking use, on the first controller module: `storage port modify -node` *`node name`* `-port` *`port name`* `-mode network` You must run this command for each port you are reconfiguring. b. Repeat the above step to reconfigure the ports on the second controller module. c. Go to substep 8c to verify all port changes.
ONTAP 9.7	a. Reconfigure the non-dedicated ports for networking use, on the first controller module: `storage port disable -node` *`node name`* `-port` *`port name`* You must run this command for each port you are reconfiguring. b. Repeat the above step to reconfigure the ports on the second controller module. c. Go to substep 8c to verify all port changes.

c. Verify that the non-dedicated ports of both controller modules are reconfigured for networking use: `storage port show`

You can enter the command on either controller module.

If your HA pair is running ONTAP 9.8 or later, the non-dedicated ports display `network` in the `Mode` column.

If your HA pair is running ONTAP 9.7, the non-dedicated ports, which display `false` in the `Is Dedicated?` column, also display `disabled` in the `State` column.

Monitor drive shelf LEDs - NS224 shelves

You can monitor the health of your drive shelf by understanding the location and status conditions of the LEDs on your drive shelf components.

- The location (blue) LEDs, on a shelf's operator display panel (ODP) and both NSMs, can be activated to

aid in physically locating the shelf that needs servicing: `storage shelf location-led modify -shelf-name shelf_name -led-status on`

If you do not know the *shelf_name* of the affected shelf, run the `storage shelf show` command.

Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the `off` option.

- An LED state can be:
  - "On": The LED illumination is solid/steady
  - "Off": The LED is not illuminated
  - "Blink": The LED turns on and off at varying intervals depending on the FRU status
  - "Any state": The LED can be "On", "Off", or "Blink"

Operator display panel LEDs

The LEDs on the drive shelf front operator display panel (ODP) indicate whether your drive shelf is functioning normally or there are problems with the hardware.

The following illustration and table describes the three LEDs on the ODP:

LED icon	LED name & color	State	Description
	Power (Green)	On	One or more power supplies are supplying power to the drive shelf.

LED icon	LED name & color	State	Description
!	Attention (Amber)	On	• An error occurred with the function of one of more shelf FRUs. Check event messages to determine corrective action to take. • If the two-digit shelf ID is also blinking, the shelf ID is in a pending state. Power cycle the drive shelf for the shelf ID to take affect.
	Location (Blue)	On	The system administrator activated this LED function.

NSM LEDs

The LEDs on an NSM indicate whether the module is functioning normally, whether it is ready for I/O traffic, and whether there are any problems with the hardware.

The following illustration and tables describe NSM LEDs associated with the function of a module and the function of each NVMe port on a module.

NSM100 modules

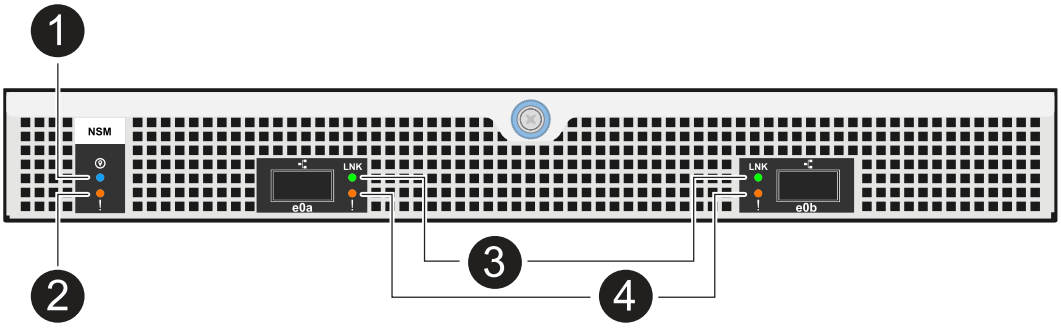

Call out	LED icon	Color	Description
1	⦿	Blue	NSM: Location
2	!	Amber	NSM: Attention
3	LNK	Green	NVMe port/link: status
4	!	Amber	NVMe port/link: attention

Status	NSM Attention (Amber)	Port LNK (Green)	Port Attention (Amber)
NSM normal	Off	Any state	Off
NSM fault	On	Any state	Any state
NSM VPD Error	On	Any state	Any state
No host port connection	Any state	Off	Off
Host port connection link active	Any state	On/Blinks with activity	Any state
Host port connection w/ fault	On	On/Off if all lanes are faulted	On
BIOS boot from BIOS image after power up	Blink	Any state	Any state

NSM100B modules

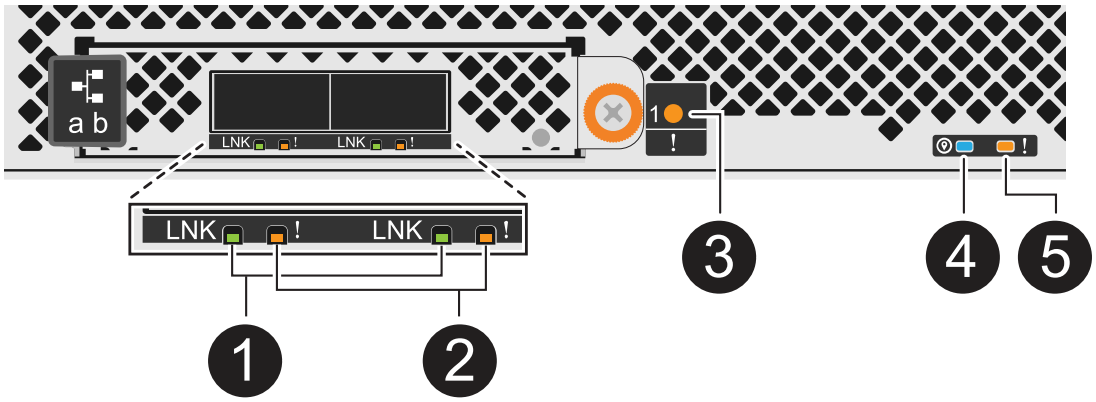

Call out	LED icon	Color	Description
1	LNK	Green	NVMe port/link: status
2	!	Amber	NVMe port/link: attention
3	!	Amber	I/O module: attention
4		Blue	NSM: Location
5	!	Amber	NSM: Attention

Status	NSM Attention (Amber)	Port LNK (Green)	Port Attention (Amber)	I/O Module Attention
NSM normal	Off	Any state	Off	Off
NSM fault	On	Any state	Any state	Off
NSM VPD Error	On	Any state	Any state	Off
No host port connection	Any state	Off	Off	Off
Host port connection link active	Any state	On/Blinks with activity	Any state	Off
Host port connection w/ fault	On	On/Off if all lanes are faulted	On	Off
BIOS boot from BIOS image after power up	Blink	Any state	Any state	Off

Status	NSM Attention (Amber)	Port LNK (Green)	Port Attention (Amber)	I/O Module Attention
I/O Module is missing	On	N/A	N/A	On

Power supply LEDs

The LEDs on an AC or DC power supply (PSU) indicate whether the PSU is functioning normally or there are hardware problems.

The following illustration and tables describe the LED on a PSU. (The illustration is an AC PSU; however, the LED location is the same on the DC PSU):

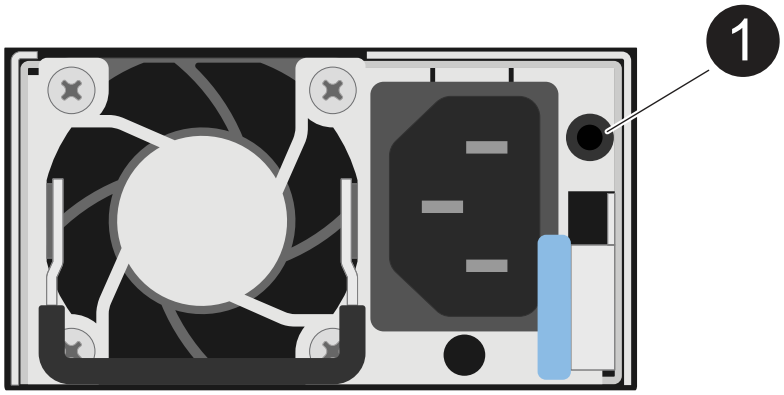

Call out	Description
1	The bi-color LED indicates power/activity when green and a fault when red.

Status	Power/activity (Green)	Attention (Red)
No AC/DC power to the enclosure	Off	Off
No AC/DC power to the PSU	Off	On
AC/DC power on, but PSU not in enclosure	Blink	Off
PSU operating correctly	On	Off
PSU failure	Off	On
Fan failure	Off	On
Firmware update mode	Blink	Off

Drive LEDs

The LEDs on an NVMe drive indicates whether it is functioning normally or there are problems with the hardware.

The following illustration and tables describe the two LEDs on an NVMe drive:

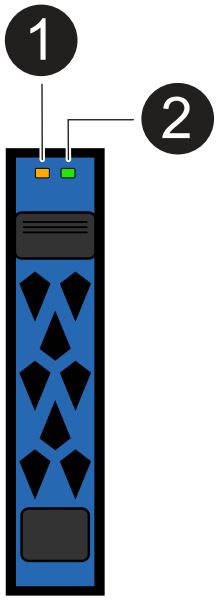

Call out	LED name	Color
1	Attention	Amber
2	Power/activity	Green

Status	Power/Activity (Green)	Attention (Amber)	Associated ODP LED
Drive installed and operational	On/Blinks with activity	Any state	N/A
Drive failure	On/Blinks with activity	On	Attention (Amber)
SES device identify set	On/Blinks with activity	Blinks	Attention (Amber) is off
SES device fault bit set	On/Blinks with activity	On	Attention (Amber)
Power control circuit failure	Off	Any state	Attention (Amber)

Replace a fan module - NS224 shelves

If one or both of the fans in your fan module fail, you can replace your fan module. This procedure can be completed nondisruptively in an NS224 drive shelf that is powered on with I/O in progress.

Before you begin

The shelf's partner NSM must be up and running, and be cabled correctly so that your shelf maintains connectivity when you remove the NSM with the failed FRU (target NSM).

NetApp Downloads: Config Advisor

About this task

- Allow at least 70 seconds between removal and installation of the NVMe shelf module (NSM).

  This allows enough time for ONTAP to process the NSM removal event.

- **Best practice:** The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.
  - To update your
    NetApp Downloads: Disk Shelf Firmware

    NetApp Downloads: Disk Drive Firmware

    Do not revert firmware to a version that does not support your shelf and its components.

- If needed, you can turn on the shelf's location (blue) LEDs to aid in physically locating the affected shelf:
  `storage shelf location-led modify -shelf-name shelf_name -led-status on`

  If you do not know the *`shelf_name`* of the affected shelf, run the `storage shelf show` command.

  A shelf has three location LEDs: one on the operator display panel and one on each NSM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the `off` option.

- When you unpack the replacement fan, save all packing materials for use when you return the failed fan.

  If you need the RMA number or additional help with the replacement procedure, contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific).

NSM100 modules

You can use the following animation to assist in replacing a fan in an NS224 with NSM100 modules.

Replace a fan in an NS224 shelf

Steps

1. Properly ground yourself.
2. Disconnect the cabling from the NSM that contains the FRU that you are replacing:
   a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling from the NSM ports.

      Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports when you reinsert the NSM later in this procedure.

3. Remove the NSM from the shelf:
   a. Loop your index fingers through the finger holes of the latching mechanisms on either side of the NSM.

      If you are removing the bottom NSM, and if the bottom rail is obstructing access to the latching mechanisms, place your index fingers through the finger holes from the inside (by crossing your arms).

   b. With your thumbs, press down and hold the orange tabs on top of the latching mechanisms.

      The latching mechanisms raise, clearing the latching pins on the shelf.

   c. Gently pull until the NSM is about one third of the way out of the shelf, grasp the NSM sides with both hands to support its weight, and then place it on a flat stable surface.

      When you begin pulling, the latching mechanism arms extend from the NSM and lock in their fully extended position.

4. Loosen the NSM cover thumb screw and open the cover.

   The FRU label on the NSM cover shows the location of the five fans, along the rear wall of the NSM.

5. Physically identify the failed fan.

   When a fan fails, the system logs a warning message to the system console indicating which fan failed.

6. Replace the failed fan:
   a. Remove the failed fan by firmly grasping the sides, where the blue touch points are located, and then lift it vertically to disconnect it from the socket.

b. Insert the replacement fan by aligning it within the guides, and then push down until the fan module connector is fully seated in the socket.

7. Close the NSM cover, and then tighten the thumb screw.
8. Reinsert the NSM into the shelf:
   a. Make sure that the latching mechanism arms are locked in the fully extended position.
   b. Using both hands, gently slide the NSM into the shelf until the weight of the NSM is fully supported by the shelf.
   c. Push the NSM into the shelf until it stops (about half an inch from the back of the shelf).

      You can place your thumbs on the orange tabs on the front of each finger loop (of the latching mechanism arms) to push in the NSM.

   d. Loop your index fingers through the finger holes of the latching mechanisms on either side of the NSM.

      If you are inserting the bottom NSM, and if the bottom rail is obstructing access to the latching mechanisms, place your index fingers through the finger holes from the inside (by crossing your arms).

   e. With your thumbs, press down and hold the orange tabs on top of the latching mechanisms.
   f. Gently push forward to get the latches over the stop.
   g. Release your thumbs from the tops of the latching mechanisms, and then continue pushing until the latching mechanisms snap into place.

      The NSM should be fully inserted into the shelf and flush with the edges of the shelf.

9. Reconnect the cabling to the NSM:
   a. Reconnect the storage cabling to the same two NSM ports.

      Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

   b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      When functioning correctly, a power supply's bicolored LED illuminates green.

      Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

10. Verify that the attention (amber) LEDs on the NSM containing the failed fan and the shelf operator display panel are no longer illuminated.

    The NSM attention LEDs turn off after the NSM reboots and no longer detects a fan issue. This can take three to five minutes.

11. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

    If any cabling errors are generated, follow the corrective actions provided.

NetApp Downloads: Config Advisor

NSM100B modules

Steps

1. Properly ground yourself.
2. Disconnect the cabling from the NSM that contains the FRU that you are replacing:
   a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling from the NSM ports.

      Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports when you reinsert the NSM later in this procedure.
3. Remove the NSM:

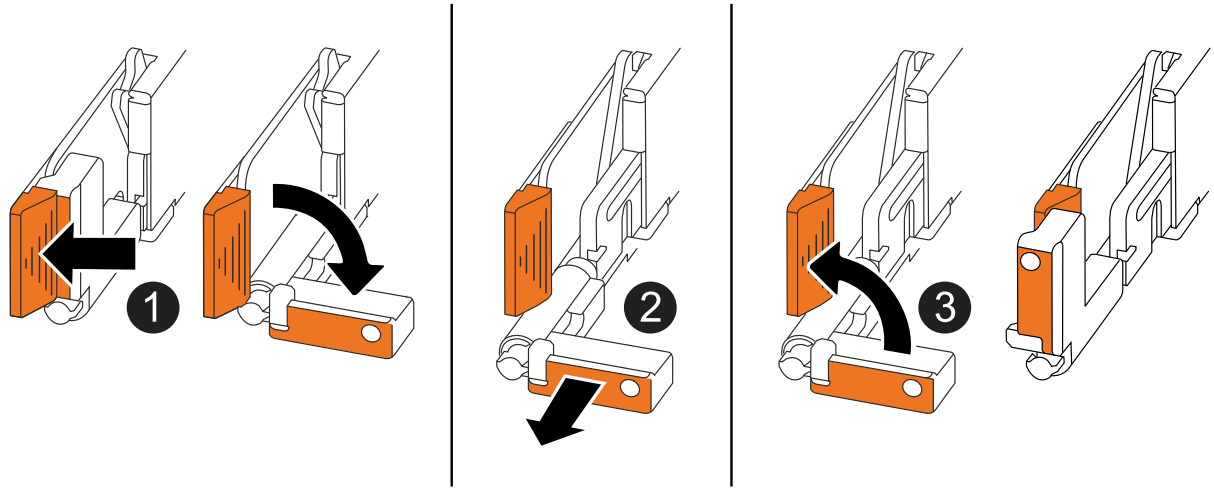

1	On both ends of the NSM, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the NSM from the midplane. As you pull, the handles extend out from the shelf. When you feel some resistance, keep pulling. • Slide the NSM out of the shelf and place it on a flat, stable surface. Make sure that you support the bottom of the NSM as you slide it out of the shelf.
3	Rotate the handles upright (next to the tabs) to move them out of the way.

4. Open the NSM cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

> The FRU label on the NSM cover shows the location of the five fans, along the rear wall of the NSM.

5. Physically identify the failed fan.

   When a fan fails, the system logs a warning message to the system console indicating which fan failed.

6. Replace the failed fan:

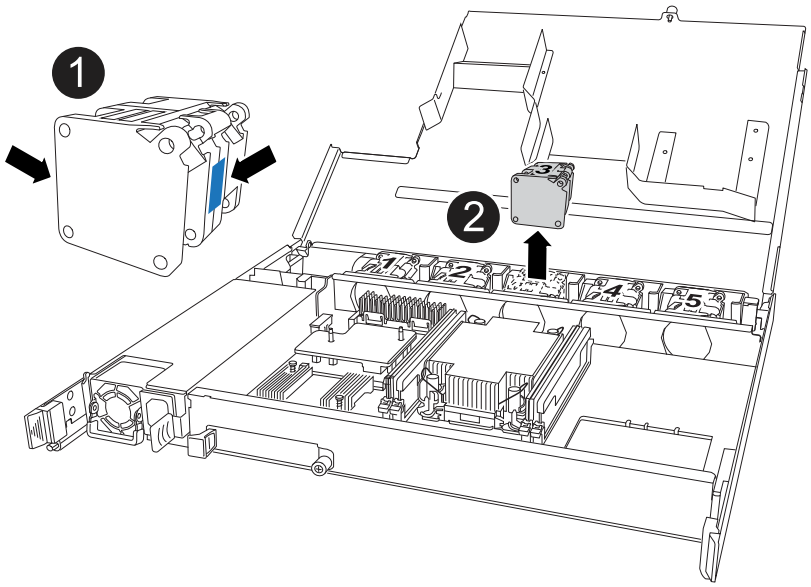

1	Remove the failed fan by firmly grasping the sides where the blue touch points are located, and then pull it straight up out of its socket.
1	Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

7. Close the NSM cover, and then tighten the thumb screw.
8. Insert the NSM into the shelf:

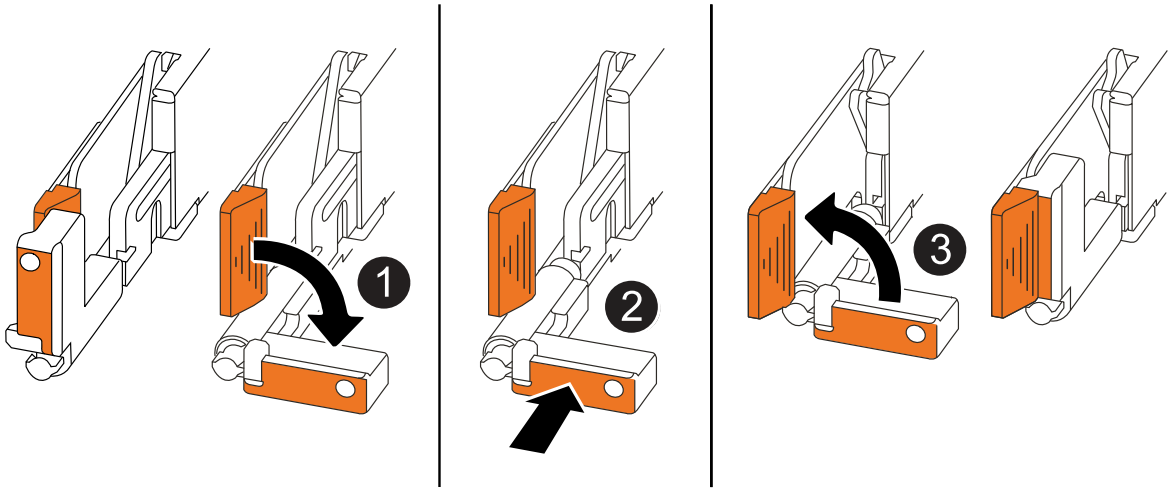

1	If you rotated the NSM handles upright (next to the tabs) to move them out of the way while you serviced the NSM, rotate them down to the horizontal position.
2	Align the rear of the NSM with the opening in the shelf, and then gently push the NSM using the handles until it is fully seated.
3	Rotate the handles to the upright position and lock in place with the tabs.

9. Reconnect the cabling to the NSM:
   a. Reconnect the storage cabling to the same two NSM ports.

      Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

   b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      When functioning correctly, a power supply's bicolored LED illuminates green.

      Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

10. Verify that the attention (amber) LEDs on the NSM containing the failed fan and the shelf operator display panel are no longer illuminated.

    The NSM attention LEDs turn off after the NSM reboots and no longer detects a fan issue. This can take three to five minutes.

11. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

    If any cabling errors are generated, follow the corrective actions provided.

    NetApp Downloads: Config Advisor

Replace the Ethernet I/O module - NS224 shelves with NSM100B modules

You can replace a failed Ethernet I/O module nondisruptively in an NS224 drive shelf that is powered on, and while I/O is in progress. This procedure only applies to NS224 shelves with NSM100B modules.

Before you begin

- The shelf's partner NSM must be up and running, and be cabled correctly so that your shelf maintains connectivity when you remove the failed NSM.

  NetApp Downloads: Config Advisor

- All other components in the system must be functioning properly.

About this task

- Allow at least 70 seconds between removal and installation of the NVMe shelf module (NSM).

  This allows enough time for ONTAP to process the NSM removal event.

- **Best practice:** The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.

  NetApp Downloads: Disk Shelf Firmware

  NetApp Downloads: Disk Drive Firmware

  > Do not revert firmware to a version that does not support your shelf and its components.

- Shelf (NSM) firmware is automatically updated (nondisruptively) on a new NSM that has a non-current firmware version.

  NSM firmware checks occur every 10 minutes. An NSM firmware update can take up to 30 minutes.

- If needed, you can turn on the shelf's location (blue) LEDs to aid in physically locating the affected shelf: `storage shelf location-led modify -shelf-name shelf_name -led-status on`

  If you do not know the `shelf_name` of the affected shelf, run the `storage shelf show` command.

  A shelf has three location LEDs: one on the operator display panel and one on each NSM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the `off` option.

- When you unpack the replacement NSM, save all packing materials for use when you return the failed NSM.

  If you need the RMA number or additional help with the replacement procedure, contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific).

Steps

1. Properly ground yourself.
2. Disconnect the cabling from the NSM that contains the FRU that you are replacing:

a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

   Power supplies do not have a power switch.

b. Disconnect the storage cabling from the NSM ports.

   Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports when you reinsert the NSM, later in this procedure.

3. Remove the NSM:

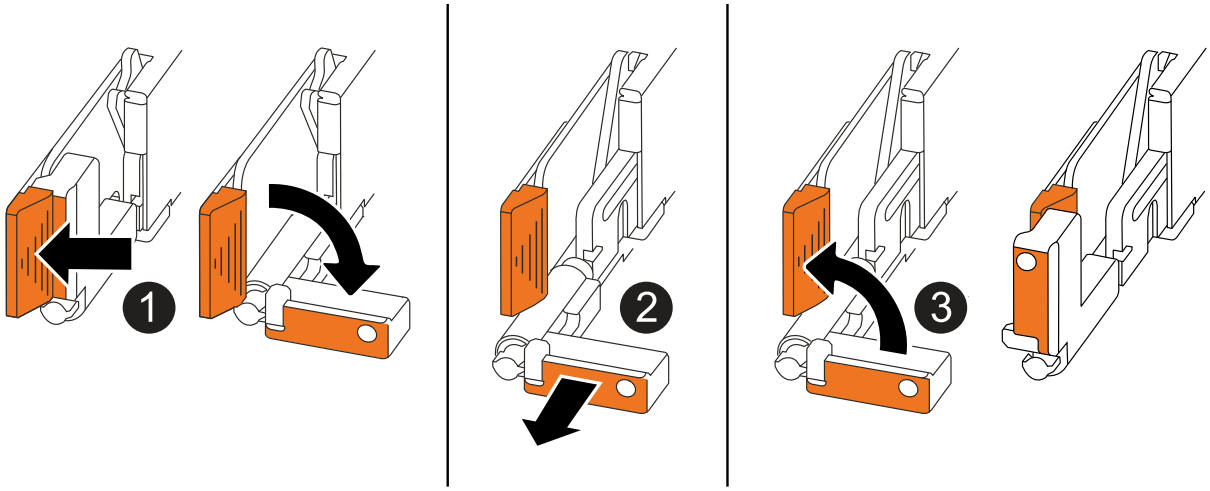

1	On both ends of the NSM, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the NSM from the midplane. As you pull, the handles extend out from the shelf. When you feel some resistance, keep pulling. • Slide the NSM out of the shelf and place it on a flat, stable surface. Make sure that you support the bottom of the NSM as you slide it out of the shelf.
3	Rotate the handles upright (next to the tabs) to move them out of the way.

4. Remove the failed I/O module from the NSM:

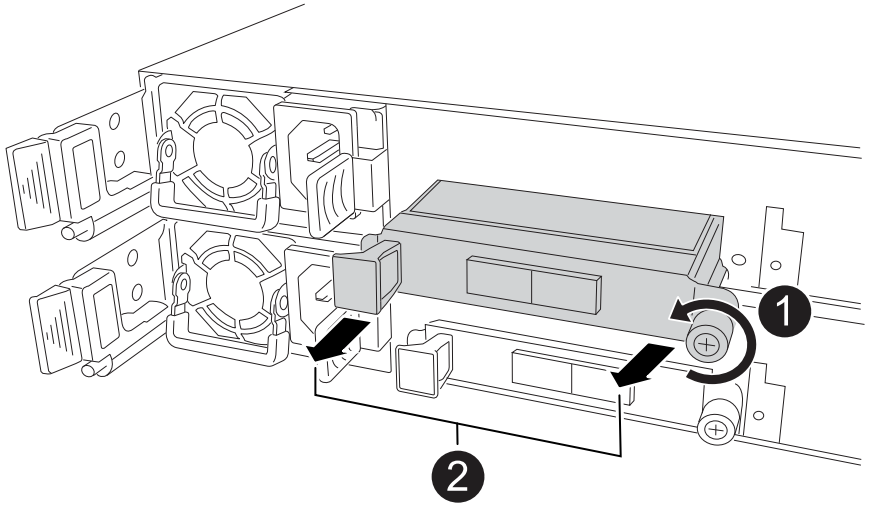

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the NSM using the port label tab on the left and the thumbscrew.

5. Install the replacement I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
6. Insert the NSM into the shelf:

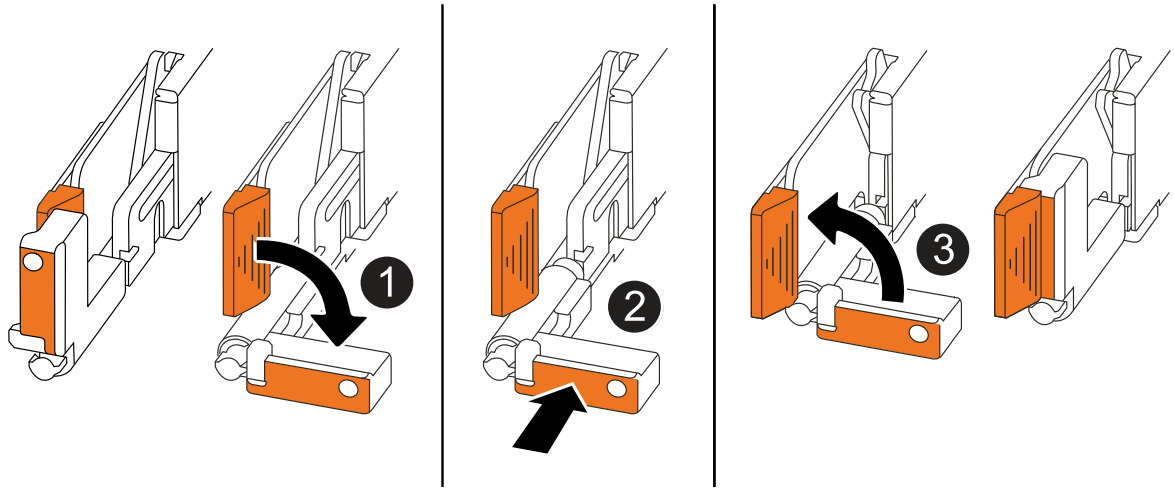

1	If you rotated the NSM handles upright (next to the tabs) to move them out of the way while you serviced the NSM, rotate them down to the horizontal position.

2	Align the rear of the NSM with the opening in the shelf, and then gently push the NSM using the handles until it is fully seated.
3	Rotate the handles to the upright position and lock in place with the tabs.

7. Recable the NSM.
   a. Reconnect the storage cabling to the same two NSM ports.

      Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

   b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      When functioning correctly, a power supply's bicolored LED illuminates green.

      Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

8. Verify that the attention (amber) LEDs on the NSM containing the failed I/O module and the shelf operator display panel are no longer illuminated

   The NSM attention LEDs turn off after the NSM reboots and no longer detects an I/O module issue. This can take three to five minutes.

9. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

   If any cabling errors are generated, follow the corrective actions provided.

   NetApp Downloads: Config Advisor

Replace an NSM - NS224 shelves

You can replace an impaired NVMe shelf module (NSM) nondisruptively in an NS224 drive shelf that is powered on, and while I/O is in progress.

Before you begin

- The shelf's partner NSM must be up and running, and be cabled correctly so that your shelf maintains connectivity when you remove the failed NSM.

  NetApp Downloads: Config Advisor

- All other components in the system must be functioning properly.

About this task

- Replacing the NSM involves moving the following:
  - **NSM100 modules**: DIMMs, fans and power supply from the impaired NSM to the replacement NSM.
  - **NSM100B modules**: DIMMs, fans, boot media, I/O module, and power supply from the impaired NSM to the replacement NSM.

You do not move the real-time clock (RTC) battery. They come preinstalled in the replacement NSM.

- Allow at least 70 seconds between removal and installation of the NVMe shelf module (NSM).

  This allows enough time for ONTAP to process the NSM removal event.

- **Best practice:** The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.

  NetApp Downloads: Disk Shelf Firmware

  NetApp Downloads: Disk Drive Firmware

  Do not revert firmware to a version that does not support your shelf and its components.

- Shelf (NSM) firmware is automatically updated (nondisruptively) on a new NSM that has a non-current firmware version.

  NSM firmware checks occur every 10 minutes. An NSM firmware update can take up to 30 minutes.

- If needed, you can turn on the shelf's location (blue) LEDs to aid in physically locating the affected shelf: `storage shelf location-led modify -shelf-name shelf_name -led-status on`

  If you do not know the `shelf_name` of the affected shelf, run the `storage shelf show` command.

  A shelf has three location LEDs: one on the operator display panel and one on each NSM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the `off` option.

- When you unpack the replacement NSM, save all packing materials for use when you return the failed NSM.

  If you need the RMA number or additional help with the replacement procedure, contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific).

NSM100 modules

You can use the following animation or the written steps to replace an NSM.

Replace an NSM in an NS224 shelf

Steps

1. Properly ground yourself.
2. Physically identify the impaired NSM.

   The system logs a warning message to the system console indicating which module is impaired. Additionally, the attention (amber) LED on the drive shelf operator display panel and the impaired module illuminate.

3. Disconnect the cabling from the impaired NSM:
   a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumbs screws if it is a DC power supply, and then unplug the power cord from the power supply.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling from the NSM ports.

      Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports on the replacement NSM, later in this procedure.

4. Remove the NSM from the shelf:
   a. Loop your index fingers through the finger holes of the latching mechanisms on either side of the NSM.

      If you are removing the bottom NSM, and if the bottom rail is obstructing access to the latching mechanisms, place your index fingers through the finger holes from the inside (by crossing your arms).

   b. With your thumbs, press down and hold the orange tabs on top of the latching mechanisms.

      The latching mechanisms raise, clearing the latching pins on the shelf.

   c. Gently pull until the NSM is about one third of the way out of the shelf, grasp the NSM sides with both hands to support its weight, and then place it on a flat stable surface.

      When you begin pulling, the latching mechanism arms extend from the NSM and lock in their fully extended position.

5. Unpack the replacement NSM, and set it on a level surface near the impaired NSM.
6. Open the cover of the impaired NSM and the replacement NSM by loosening the thumbscrew on each cover.

   The FRU label on the NSM cover shows the location of the DIMMs and fans.

7. Move the DIMMs from the impaired NSM to the replacement NSM:

a. Note the orientation of the DIMMs in the slots so that you can insert the DIMMs into the replacement NSM using the same orientation.

b. Eject a DIMM from its slot by slowly pushing apart the ejector tabs at both ends of the DIMM slot, and then lift the DIMM out of the slot.

> Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
> The ejector tabs remain in the open position.

c. Hold the DIMM by the corners, and then insert the DIMM squarely into a slot on the replacement NSM.

   The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

   When inserted correctly, the DIMM should go in easily but fit tightly in the slot. If not, reinsert the DIMM.

d. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

e. Repeat substeps 7a through 7d for the remaining DIMMs.

8. Move the fans from the impaired NSM to the replacement NSM:

   a. Firmly grasp a fan from the sides, where the blue touch points are located, and then lift it vertically to disconnect it from the socket.

      You might need to gently rock the fan back and forth to disconnect it before lifting it out.

   b. Align the fan with the guides in the replacement NSM, and then push down until the fan module connector is fully seated in the socket.

   c. Repeat substeps 8a and 8b for the remaining fans.

9. Close the cover of each NSM, and then tighten each thumbscrew.

10. Move the power supply from the impaired NSM to the replacement NSM:

    a. Rotate the handle up, to its horizontal position, and then grasp it.

    b. With your thumb, press the blue tab to release the locking mechanism.

    c. Pull the power supply out of the NSM while using your other hand to support its weight.

    d. Using both hands, support and align the edges of the power supply with the opening in the replacement NSM.

    e. Gently push the power supply into the NSM until the locking mechanism clicks into place.

    > Do not use excessive force or you might damage the internal connector.

    f. Rotate the handle down, so it is out of the way of normal operations.

11. Insert the replacement NSM into the shelf:

    a. Make sure that the latching mechanism arms are locked in the fully extended position.

    b. Using both hands, gently slide the NSM into the shelf until the weight of the NSM is fully supported by the shelf.

    c. Push the NSM into the shelf until it stops (about half an inch from the back of the shelf).

You can place your thumbs on the orange tabs on the front of each finger loop (of the latching mechanism arms) to push in the NSM.

d. Loop your index fingers through the finger holes of the latching mechanisms on either side of the NSM.

> If you are inserting the bottom NSM, and if the bottom rail is obstructing access to the latching mechanisms, place your index fingers through the finger holes from the inside (by crossing your arms).

e. With your thumbs, press down and hold the orange tabs on top of the latching mechanisms.

f. Gently push forward to get the latches over the stop.

g. Release your thumbs from the tops of the latching mechanisms, and then continue pushing until the latching mechanisms snap into place.

The NSM should be fully inserted into the shelf and flush with the edges of the shelf.

12. Reconnect the cabling to the NSM:

a. Reconnect the storage cabling to the same two NSM ports.

Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power supply.

When functioning correctly, a power supply's bicolored LED illuminates green.

Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

13. Verify that the attention (amber) LED on the shelf operator display panel is no longer illuminated.

The operator display panel attention LED turns off after the NSM reboots. This can take three to five minutes.

14. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

If any cabling errors are generated, follow the corrective actions provided.

NetApp Downloads: Config Advisor

15. Make sure that both NSMs in the shelf are running the same version of firmware: version 0200 or later.

NSM100B modules

Steps

1. Properly ground yourself.
2. Physically identify the impaired NSM.

The system logs a warning message to the system console indicating which module is impaired. Additionally, the attention (amber) LED on the drive shelf operator display panel and the impaired

module illuminate.

3. Disconnect the cabling from the impaired NSM:
   a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumbs screws if it is a DC power supply, and then unplug the power cord from the power supply.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling from the NSM ports.

      Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports on the replacement NSM, later in this procedure.

4. Remove the NSM:

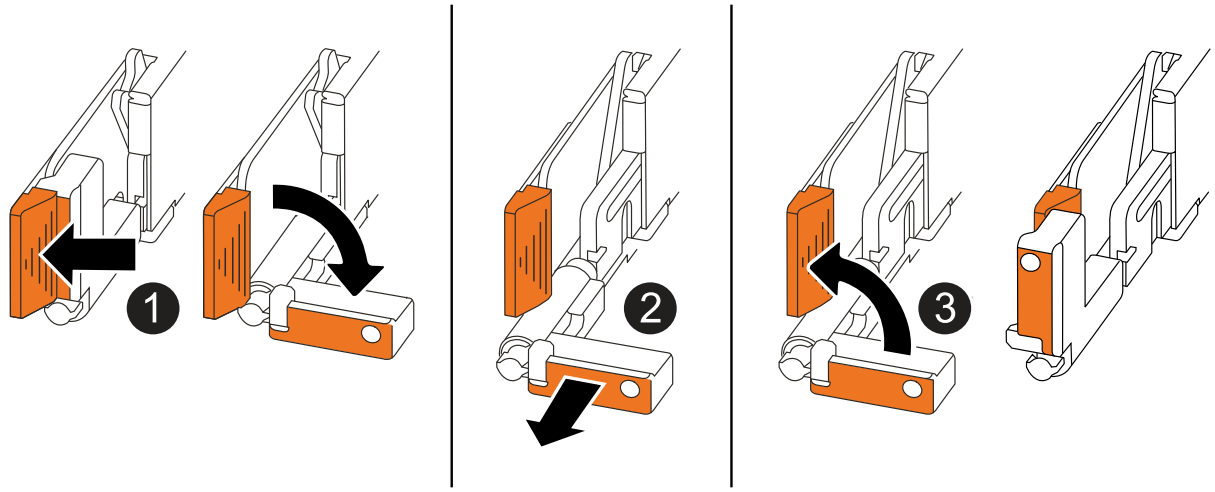

1	On both ends of the NSM, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the NSM from the midplane. As you pull, the handles extend out from the shelf. When you feel some resistance, keep pulling. • Slide the NSM out of the shelf and place it on a flat, stable surface. Make sure that you support the bottom of the NSM as you slide it out of the shelf.
3	Rotate the handles upright (next to the tabs) to move them out of the way.

5. Unpack the replacement NSM, and set it on a level surface near the impaired NSM.
6. Open the covers of both NSMs by loosening the thumbscrew on each cover.
7. Move the DIMMs from the impaired NSM to the replacement to the replacement NSM:
   a. Remove DIMM from the impaired NSM:

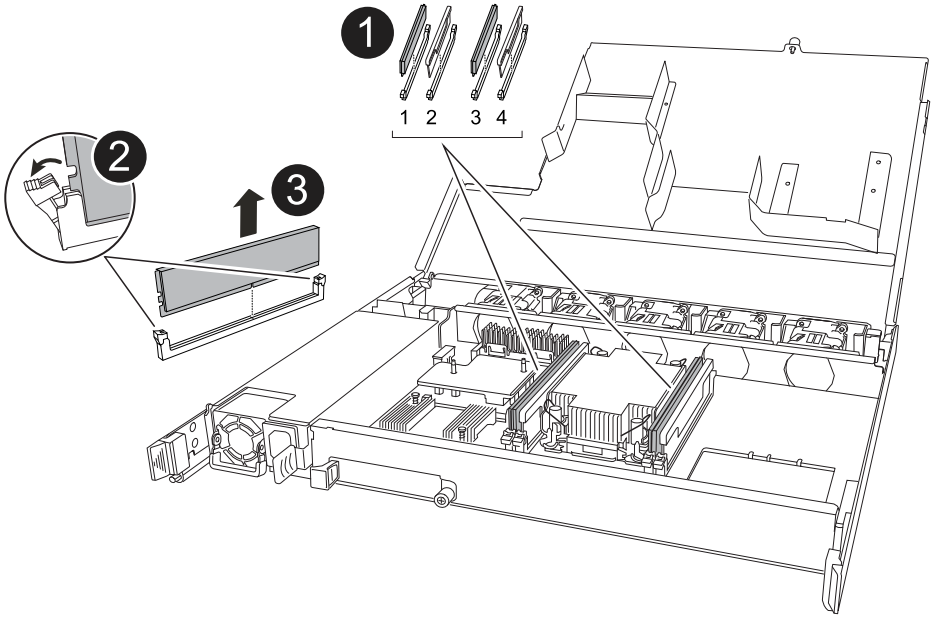

1	DIMM slot numbering and positions. The NSM contains DIMMs in slots 1 and 3, and DIMM Blanks in slots 2 and 4.
2	• Note the orientation of the DIMM in the socket so that you can insert it into the replacement DIMM using the same orientation. • Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

b. Install DIMM in the replacement NSM:

i. Hold the DIMM by the corners, and then insert the DIMM squarely into a slot.

The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

When inserted correctly, the DIMM should go in easily but fit tightly in the slot. If not, reinsert the DIMM.

ii. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

iii. Repeat for the other DIMM.

8. Move all fans from the impaired NSM to the replacement NSM:

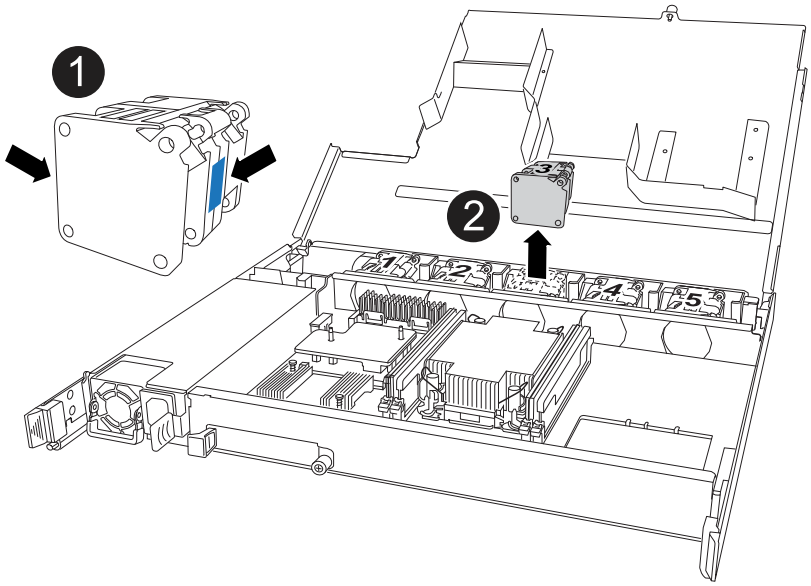

1	Remove the failed fan by firmly grasping the sides where the blue touch points are located, and then pull it straight up out of its socket.
1	Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

9. Move the boot media to the replacement NSM:
   a. Remove the boot media from the impaired NSM:

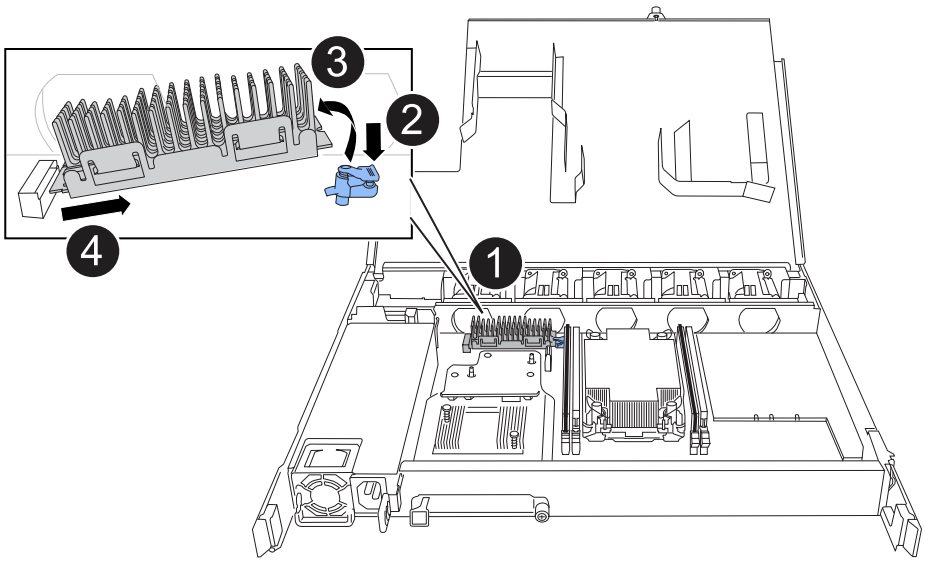

1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.

3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

b. Install the boot media in the replacement NSM:

i. Align the edges of the boot media with the socket housing in the replacement NSM, and then gently push it squarely into the socket.

ii. Rotate the boot media down toward the locking button.

iii. Push the locking button, rotate the boot media all the way down, and then release the locking button.

10. Move the I/O module from the impaired NSM to the replacement NSM.

a. Remove I/O module from the impaired NSM:

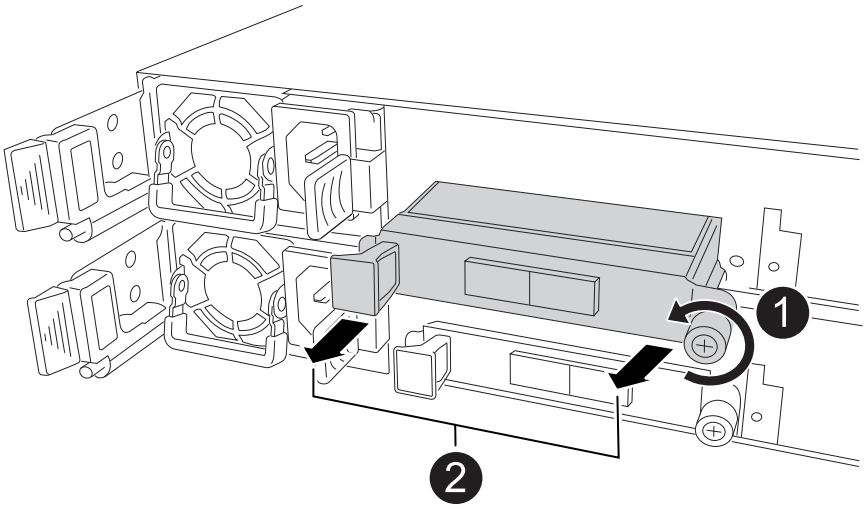

1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the NSM using the port label tab on the left and the thumbscrew.

b. Install I/O module in the replacement NSM:

i. Align the I/O module with the edges of the slot in the replacement NSM.

ii. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O module.

11. Close the cover of each NSM, and then tighten each thumbscrew.

12. Move the power supply from the impaired NSM to the replacement NSM:

a. Rotate the handle up, to its horizontal position, and then grasp it.

b. With your thumb, press the blue tab (AC PSU) or terra cotta tab (DC PSU) to release the locking

mechanism.

c. Pull the power supply out of the NSM while using your other hand to support its weight.

d. Using both hands, support and align the edges of the power supply with the opening in the replacement NSM.

e. Gently push the power supply into the NSM until the locking mechanism clicks into place.

> Do not use excessive force or you might damage the internal connector.

f. Rotate the handle down, so it is out of the way of normal operations.

13. Insert the NSM into the shelf:

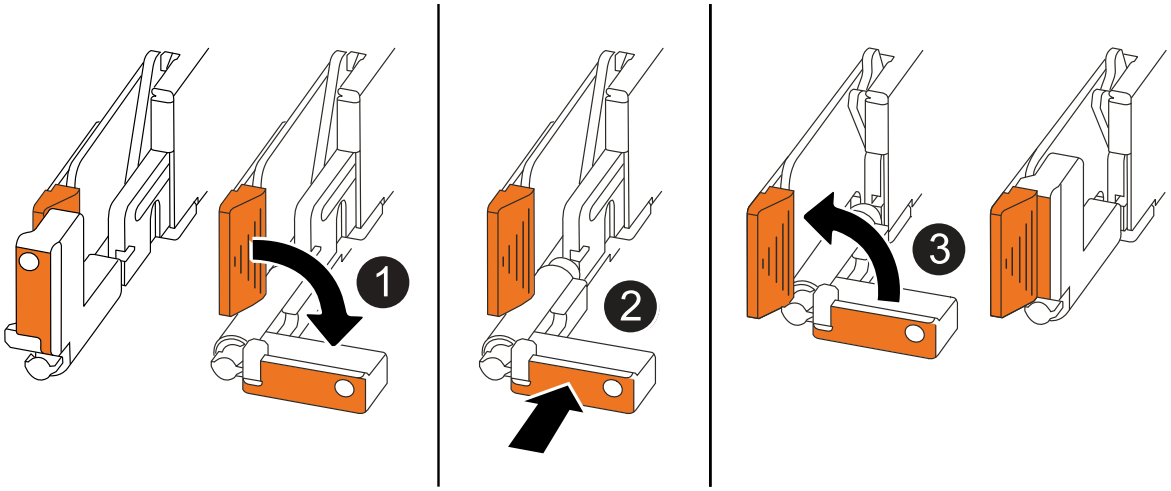

1	If you rotated the NSM handles upright (next to the tabs) to move them out of the way while you serviced the NSM, rotate them down to the horizontal position.
2	Align the rear of the NSM with the opening in the shelf, and then gently push the NSM using the handles until it is fully seated.
3	Rotate the handles to the upright position and lock in place with the tabs.

14. Reconnect the cabling to the NSM:

a. Reconnect the storage cabling to the same two NSM ports.

Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power supply.

When functioning correctly, a power supply's bicolored LED illuminates green.

Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

15. Verify that the attention (amber) LED on the shelf operator display panel is no longer illuminated.

    The operator display panel attention LED turns off after the NSM reboots. This can take three to five minutes.

16. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

    If any cabling errors are generated, follow the corrective actions provided.

    NetApp Downloads: Config Advisor

17. Make sure that both NSMs in the shelf are running the same version of firmware: version 0300 or later.

Hot-swap a power supply - NS224 shelves

You can replace a failed power supply nondisruptively in an NS224 drive shelf that is powered on, and while I/O is in progress.

About this task

- This procedure applies to NS224 shelves with either NSM100 or NSM100B modules.
- Do not mix power supplies with different efficiency ratings or with different input types.

  Always replace like for like.

- If you are replacing more than one power supply, you must do so one at a time so that the shelf maintains power.
- **Best practice:** The best practice is to replace the power supply within two minutes of removal from the NSM.

  If you exceed the two minutes, the shelf continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- Do not revert firmware to a version that does not support your shelf and its components.
- If needed, you can turn on the shelf's location (blue) LEDs to aid in physically locating the affected shelf: `storage shelf location-led modify -shelf-name` *`shelf_name`* `-led-status on`

  If you do not know the `shelf_name` of the affected shelf, run the `storage shelf show` command.

  A shelf has three location LEDs: one on the operator display panel and one on each NSM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the off option.

- When you unpack the replacement power supply, save all packing materials for use when you return the failed power supply.

  If you need the RMA number or additional help with the replacement procedure, contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific).

- If you are replacing an AC power supply for an NS224 shelf, you can watch the following animation to familiarize yourself with the procedure before beginning. The animation shows an NS224 with NSM100

modules, but the steps are the same for shelves with NSM100B modules.

Replace AC power supply for NS224

Hot-swap an AC power supply in an NS224 shelf

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Replace an AC power supply

Complete the following steps to replace an AC power supply.

Steps

1. Properly ground yourself.
2. Physically identify the failed power supply.

   The system logs a warning message to the system console indicating which power supply failed. Additionally, the attention (amber) LED on the shelf operator display panel illuminates and the bicolored LED on the failed power supply illuminates red.

3. Disconnect the power cord from the power supply by opening the power cord retainer, and then unplug the power cord from the power supply.

   Power supplies do not have a power switch.

4. Remove the power supply:

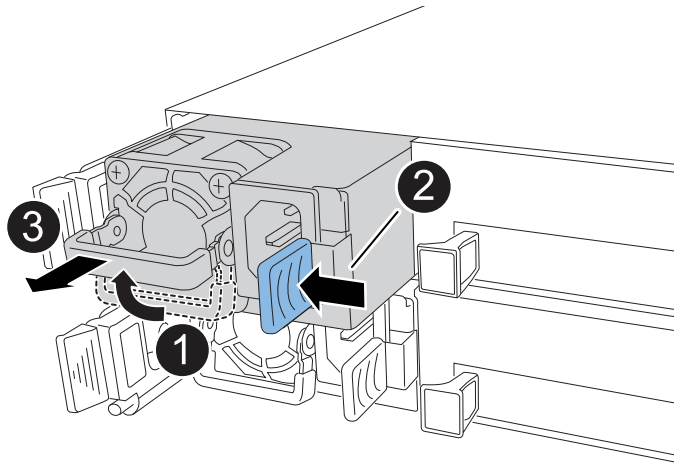

1	Rotate the handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the locking mechanism.
3	Pull the power supply out of the NSM while using your other hand to support its weight.

5. Insert the replacement power supply:
   a. Using both hands, support and align the edges of the power supply with the opening in the NSM.
   b. Gently push the power supply into the NSM until the locking mechanism clicks into place.

      Do not use excessive force or you might damage the internal connector.

   c. Rotate the handle down, so it is out of the way of normal operations.
6. Connect the power cord to the power supply and secure the power cord with the power cord retainer.

When functioning correctly, a power supply's bicolored LED illuminates green.

Option 2: Replace a DC power supply

To replace a DC power supply, complete the following steps.

1. Properly ground yourself.
2. Physically identify the failed power supply.

   The system logs a warning message to the system console indicating which power supply failed. Additionally, the attention (amber) LED on the shelf operator display panel illuminates and the bicolored LED on the failed power supply illuminates red.

3. Disconnect the power supply:
   a. Unscrew the two thumb screws on the D-SUB DC power cable connector.

      The illustration and table in step 4 shows the two thumb screws (item #1) and the D-SUB DC power cable connector (item #2).

   b. Unplug the D-SUB DC power cable connector from the power supply and set it aside.
4. Remove the power supply:
   a. Rotate the handle up, to its horizontal position, and then grasp it.
   b. With your thumb, press the terra-cotta tab to release the locking mechanism.
   c. Pull the power supply out of the NSM while using your other hand to support its weight.

   The power supply is short. Always use two hands to support it when removing it from the NSM so that it does not swing free from the NSM and injure you.

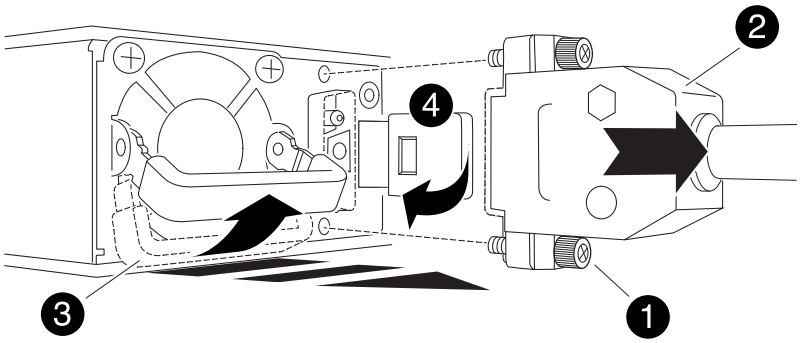

1	Thumb screws
2	D-SUB DC power cable connector
3	Power supply handle
4	Blue/Terra cotta power supply locking tab

5. Insert the replacement power supply:
   a. Using both hands, support and align the edges of the power supply with the opening in the NSM.
   b. Gently push the power supply into the NSM until the locking mechanism clicks into place.

      A power supply must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the power supply is not properly seated.

      Do not use excessive force or you might damage the internal connector.

   c. Rotate the handle down, so it is out of the way of normal operations.
6. Reconnect the D-SUB DC power cable:

   Once power is restored to the power supply, the status LED should be green.

   a. Plug the D-SUB DC power cable connector into the power supply.
   b. Tighten the two thumb screws to secure the D-SUB DC power cable connector to the power supply.

Replace the real-time clock battery - NS224 shelves

You can replace a failed real-time clock (RTC) battery nondisruptively in an NS224 drive shelf that is powered on, and while I/O is in progress.

Before you begin

- The shelf's partner NSM must be up and running, and be cabled correctly so that your shelf maintains connectivity when you remove the NSM with the failed FRU (target NSM).

  NetApp Downloads: Config Advisor

- All other components in the system must be functioning properly.

About this task

- Allow at least 70 seconds between removal and installation of the NVMe shelf module (NSM).

  This allows enough time for ONTAP to process the NSM removal event.

- After you replace the RTC battery, reinstall the NSM, and the module boots, the real-time clock time is updated by ONTAP.
- **Best practice:** The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.

  NetApp Downloads: Disk Shelf Firmware

  NetApp Downloads: Disk Drive Firmware

  Do not revert firmware to a version that does not support your shelf and its components.

- If needed, you can turn on the shelf's location (blue) LEDs to aid in physically locating the affected shelf:
  `storage shelf location-led modify -shelf-name shelf_name -led-status on`

If you do not know the `shelf_name` of the affected shelf, run the `storage shelf show` command.

A shelf has three location LEDs: one on the operator display panel and one on each NSM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the `off` option.

- When you unpack the replacement RTC battery, save all packing materials for use when you return the failed RTC battery.

  If you need the RMA number or additional help with the replacement procedure, contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific).

NSM100 modules

For NSM100 modules, you can use the following animation or the written steps to replace an RTC battery.

Replace an RTC battery in an NS224 shelf

Steps

1. Properly ground yourself.
2. Disconnect the cabling from the NSM that contains the FRU that you are replacing:
   a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling from the NSM ports.

      Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports when you reinsert the NSM, later in this procedure.

3. Remove the NSM from the shelf:
   a. Loop your index fingers through the finger holes of the latching mechanisms on either side of the NSM.

      If you are removing the bottom NSM, and if the bottom rail is obstructing access to the latching mechanisms, place your index fingers through the finger holes from the inside (by crossing your arms).

   b. With your thumbs, press down and hold the orange tabs on top of the latching mechanisms.

      The latching mechanisms raise, clearing the latching pins on the shelf.

   c. Gently pull until the NSM is about one third of the way out of the shelf, grasp the NSM sides with both hands to support its weight, and then place it on a flat stable surface.

      When you begin pulling, the latching mechanism arms extend from the NSM and lock in their fully extended position.

4. Loosen the NSM cover thumb screw and open the cover.

   The FRU label on the NSM cover shows the location of the RTC battery, near the front of the NSM and to the right of the power supply.

5. Physically identify the failed RTC battery.
6. Replace the RTC battery:
   a. Remove the battery by gently pushing it away from the holder until it is at an inclined angle (tilted away from the holder), and then lift it out of the holder.
   b. Insert the replacement battery into the holder at an inclined angle (tilted away from the holder), push it into an upright position, and then press it firmly into the connector until it is fully seated.

The positive side of the battery, marked with a plus sign, is oriented outward (away from the holder), corresponding to the plus sign marked on the NSM board.

7. Close the NSM cover, and then tighten the thumb screw.
8. Make sure that the latching mechanism arms are locked in the fully extended position.
9. Using both hands, gently slide the NSM into the shelf until the weight of the NSM is fully supported by the shelf.
10. Push the NSM into the shelf until it stops (about half an inch from the back of the shelf).

    You can place your thumbs on the orange tabs on the front of each finger loop (of the latching mechanism arms) to push in the NSM.

11. Loop your index fingers through the finger holes of the latching mechanisms on either side of the NSM.

    If you are inserting the bottom NSM, and if the bottom rail is obstructing access to the latching mechanisms, place your index fingers through the finger holes from the inside (by crossing your arms).

12. With your thumbs, press down and hold the orange tabs on top of the latching mechanisms.
13. Gently push forward to get the latches over the stop.
14. Release your thumbs from the tops of the latching mechanisms, and then continue pushing until the latching mechanisms snap into place.

    The NSM should be fully inserted into the shelf and flush with the edges of the shelf.

15. Reconnect the cabling to the NSM:
    a. Reconnect the storage cabling to the same two NSM ports.

       Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

    b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

       When functioning correctly, a power supply's bicolored LED illuminates green.

       Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

16. Verify that the attention (amber) LEDs on the NSM containing the failed RTC battery and the shelf operator display panel are no longer illuminated

    The NSM attention LEDs turn off after the NSM reboots and no longer detects an RTC battery issue. This can take three to five minutes.

17. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

    If any cabling errors are generated, follow the corrective actions provided.

NetApp Downloads: Config Advisor

NSM100B modules

Steps

1. Properly ground yourself.
2. Disconnect the cabling from the NSM that contains the FRU that you are replacing:
   a. Disconnect the power cord from the power supply by opening the power cord retainer if it is an AC power supply, or unscrewing the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      Power supplies do not have a power switch.

   b. Disconnect the storage cabling from the NSM ports.

      Make a note of the NSM ports that each cable is connected to. You reconnect the cables to the same ports when you reinsert the NSM, later in this procedure.
3. Remove the NSM:

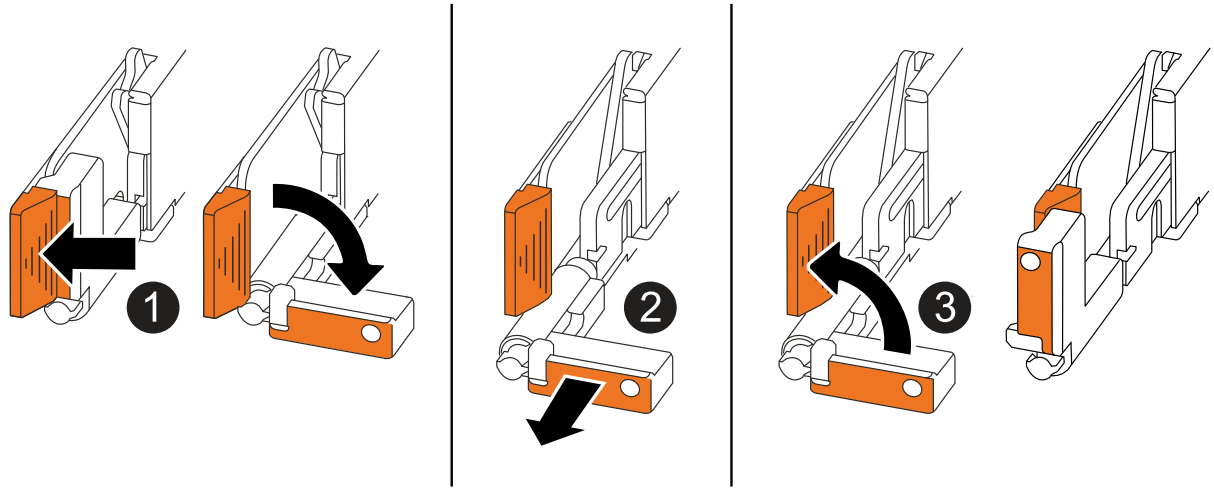

1	On both ends of the NSM, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the NSM from the midplane. As you pull, the handles extend out from the shelf. When you feel some resistance, keep pulling. • Slide the NSM out of the shelf and place it on a flat, stable surface. Make sure that you support the bottom of the NSM as you slide it out of the shelf.
3	Rotate the handles upright (next to the tabs) to move them out of the way.

4. Open the module cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.
5. Locate the RTC battery and replace it.
   a. Remove the failed battery:

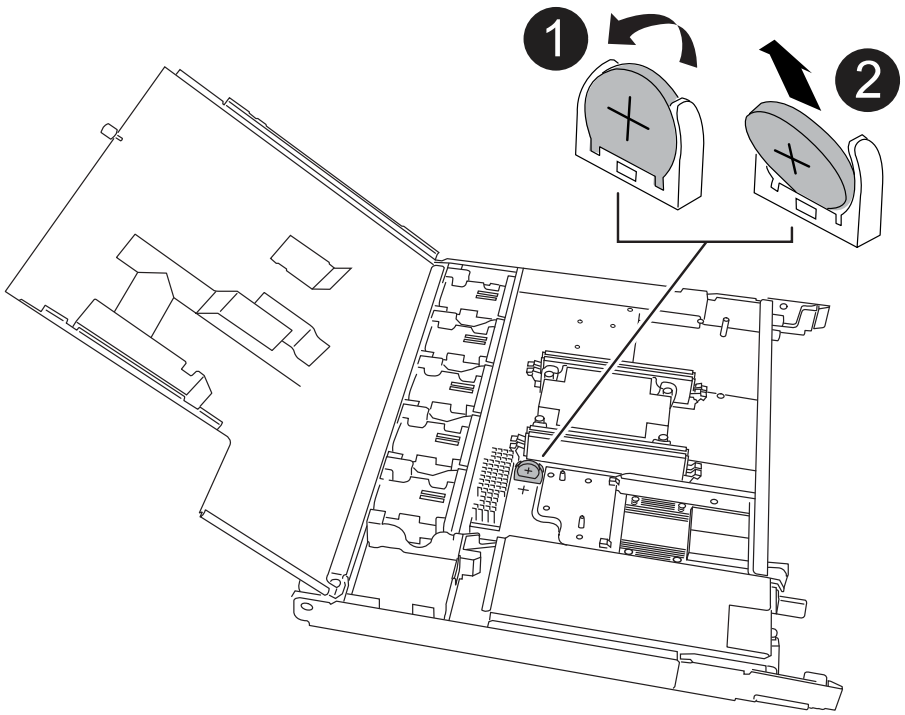

1	Gently rotate the RTC battery at an angle away from its holder.
2	Lift the RTC battery out of its holder.

   b. Remove the replacement battery from the antistatic shipping bag.
   c. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.

> You must ensure that the plus sign on the battery corresponds to the plus sign on the motherboard.

   d. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
6. Close the NSM cover and turn the thumbscrew clockwise until tightened.
7. Insert the NSM into the shelf:

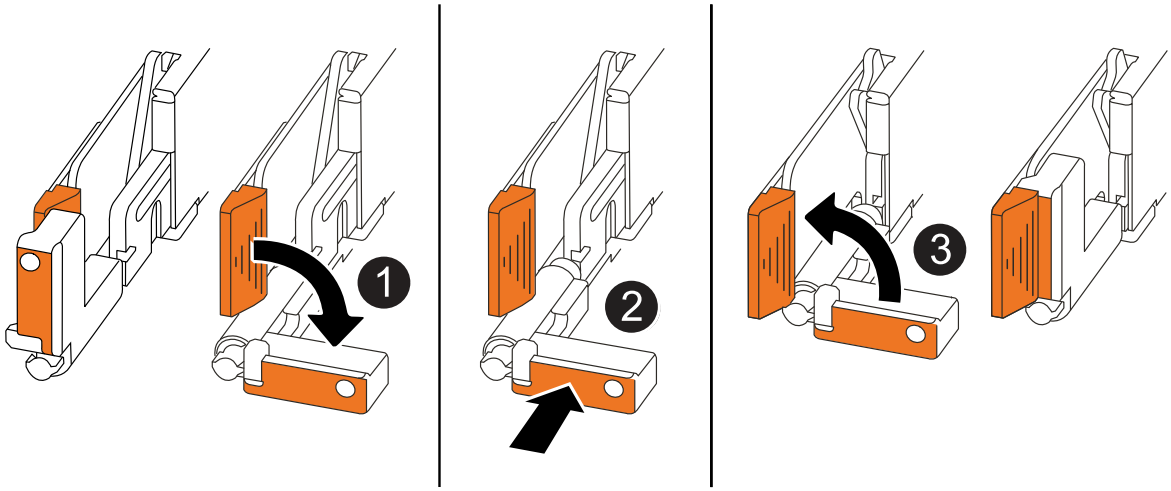

1	If you rotated the NSM handles upright (next to the tabs) to move them out of the way while you serviced the NSM, rotate them down to the horizontal position.
2	Align the rear of the NSM with the opening in the shelf, and then gently push the NSM using the handles until it is fully seated.
3	Rotate the handles to the upright position and lock in place with the tabs.

8. Recable the NSM.
   a. Reconnect the storage cabling to the same two NSM ports.

      Cables are inserted with the connector pull-tab facing up. When a cable is inserted correctly, it clicks into place.

   b. Reconnect the power cord to the power supply, and then secure the power cord with the power cord retainer if it is an AC power supply, or tighten the two thumb screws if it is a DC power supply, and then unplug the power cord from the power supply.

      When functioning correctly, a power supply's bicolored LED illuminates green.

      Additionally, both NSM port LNK (green) LEDs illuminate. If a LNK LED does not illuminate, reseat the cable.

9. Verify that the attention (amber) LEDs on the NSM containing the failed RTC battery and the shelf operator display panel are no longer illuminated

   The NSM attention LEDs turn off after the NSM reboots and no longer detects an RTC battery issue. This can take three to five minutes.

10. Verify that the NSM is cabled correctly, by running Active IQ Config Advisor.

    If any cabling errors are generated, follow the corrective actions provided.

    NetApp Downloads: Config Advisor

SAS shelves

Install and cable

Install and cable shelves - DS212C, DS224C, or DS460C

If your new system—HA pair or single-controller configuration—did not come installed in a cabinet, you can install and cable the disk shelves in a rack.

About this task

- Disk shelves with IOM12/IOM12B modules are shipped with shelf IDs preset to 00.

  If you have an HA pair with at least two stacks, the disk shelf containing the root aggregates for the second stack has the shelf ID preset to 10.

  You must set shelf IDs so they are unique within the HA pair or single-controller configuration. You can manually set shelf IDs or have shelf IDs automatically assigned for all disk shelves in the HA pair or single-controller configuration using a command in maintenance mode. Instructions for both methods are provided.

- You can identify the disk shelves containing the root aggregates by the labels on the disk shelf box and disk shelf chassis.

  The labels show the stack number; for example, **Loop or Stack #: 1** and **Loop or Stack #: 2**. Disk shelves that do not contain the root aggregates only show the disk shelf serial number on the labels.

- If at system setup and configuration, you do not configure the system to use automatic disk ownership assignment. You need to manually assign disk ownership.
- In-band Alternate Control Path (IBACP) is automatically enabled.

  IBACP is not supported on single-path HA or single-path configurations.

Before you begin

You must meet certain requirements and familiarize yourself with the best practices and considerations for this procedure before installing and cabling the disk shelves.

- Obtain the installation and setup instructions for your platform model.

  The installation and setup instructions address the complete procedure for your system installation, setup, and configuration. You only use this procedure in conjunction with the platform installation and setup instructions if you need detailed information about installing or cabling the disk shelves to your storage system.

  Installation and setup instructions can be found by navigating to your platform in the AFF and FAS System Documentation.

- Disk shelves and controllers must not be powered on at this time.
- **Best practice**: Ensure your system can recognize and utilize newly qualified disk drives by downloading the current version of the Disk Qualification Package (DQP).

  This allows you to avoid system event messages about having non-current disk drive information. You also avoid the possible prevention of disk partitioning because disk drives are not recognized. The DQP notifies

you of non-current disk drive firmware.

- **Best practice**: Verify SAS connections are cabled correctly and that shelf IDS are unique within the HA pair or single-controller configuration by downloading and running Config Advisor after a new system installation.

  If any SAS cabling or duplicate shelf ID errors are generated, follow the corrective actions provided.

  You need network access to download Config Advisor.

- Familiarize yourself with the considerations for properly handling SAS cables:
  - If you are using mini-SAS HD SAS optical cables, you must have met the rules in Mini-SAS HD SAS optical cable rules.
  - Visually inspect the SAS port to verify the proper orientation of the connector before plugging it in.

    The SAS cable connectors are keyed. When oriented correctly into a SAS port, the connector clicks into place and if the disk shelf power is on at the time, the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector).

    For controllers, the orientation of SAS ports can vary depending on the platform model; therefore, the correct orientation of the SAS cable connector varies.

  - To prevent degraded performance, do not twist, fold, pinch, or step on the cables.

    Cables have a minimum bend radius. Cable manufacturer specifications define the minimum bend radius; however, a general guideline for minimum bend radius is 10 times the cable diameter.

  - **Best practice**: Use Velcro wraps instead of tie-wraps to bundle and secure system cables to allow for easier cable adjustments.
- Familiarize yourself with the considerations for properly handling DS460C drives:
  - The drives are packaged separately from the shelf chassis.

    You should take inventory of the drives along with the rest of the system equipment you received.

  - After you unpack the drives, you should save the packaging materials for future use.

    Possible loss of data access: If in the future, you move the shelf to a different part of the data center or transport the shelf to a different location, you need to remove the drives from the drive drawers to avoid possible damage to the drive drawers and drives.

    Keep disk drives in their ESD bag until you are ready to install them.

  - When handling the drives, always wear an ESD wrist strap grounded to an unpainted surface on your storage enclosure chassis to prevent static discharges.

    If a wrist strap is unavailable, touch an unpainted surface on your storage enclosure chassis before handling the disk drive.

Step 1: Install disk shelves for a new system installation

You install the disk shelves in a rack using the rack mount kits that came with the disk shelves.

1. Install the rack mount kit (for two-post or four-post rack installations) that came with your disk shelf using the installation flyer that came with the kit.

   If you are installing multiple disk shelves, you should install them from the bottom to the top of the rack for the best stability.

   Do not flange-mount the disk shelf into a telco-type rack; the disk shelf's weight can cause it to collapse in the rack under its own weight.

2. Install and secure the disk shelf onto the support brackets and rack using the installation flyer that came with the kit.

   To make a disk shelf lighter and easier to maneuver, remove the power supplies and I/O modules (IOMs).

   Although the drives in DS460C shelves are packaged separately, which makes the shelf lighter, an empty DS460C shelf still weighs approximately 132 lb (60kg). It is recommended that you use a mechanized lift or four people using the lift handles to safely move an empty DS460C shelf.

   Your DS460C shipment includes four detachable lift handles (two for each side). To use the lift handles, install them by inserting the tabs of the handles into the slots in the side of the shelf and pushing up until they click into place. Then, as you slide the disk shelf onto the rails, detach one set of handles at a time using the thumb latch. The following illustration shows how to attach a lift handle.

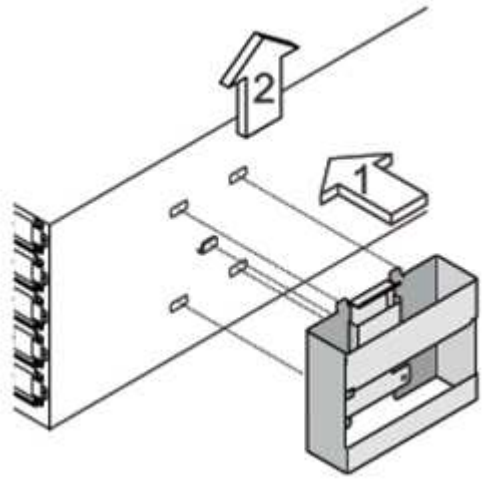

3. Reinstall any power supplies and IOMs you removed prior to installing your disk shelf into the rack.
4. If you are installing a DS460C disk shelf, install the drives into the drive drawers. Otherwise, go to the next step.

   Always wear an ESD wrist strap grounded to an unpainted surface on your storage enclosure chassis to prevent static discharges.

   If a wrist strap is unavailable, touch an unpainted surface on your storage enclosure chassis before handling the disk drive.

   If you purchased a partially populated shelf, meaning that the shelf has less than the 60 drives it supports, install the drives in each drawer as follows:

◦ Install the first four drives into the front slots (0, 3, 6, and 9).

> **Risk of equipment malfunction:** To allow for proper air flow and prevent overheating, always install the first four drives into the front slots (0, 3, 6, and 9).

◦ For the remaining drives, evenly distribute them across each drawer.

The following illustration shows how the drives are numbered from 0 to 11 in each drive drawer within the shelf.

a. Open the top drawer of the shelf.
b. Remove a drive from its ESD bag.
c. Raise the cam handle on the drive to vertical.
d. Align the two raised buttons on each side of the drive carrier with the matching gap in the drive channel on the drive drawer.

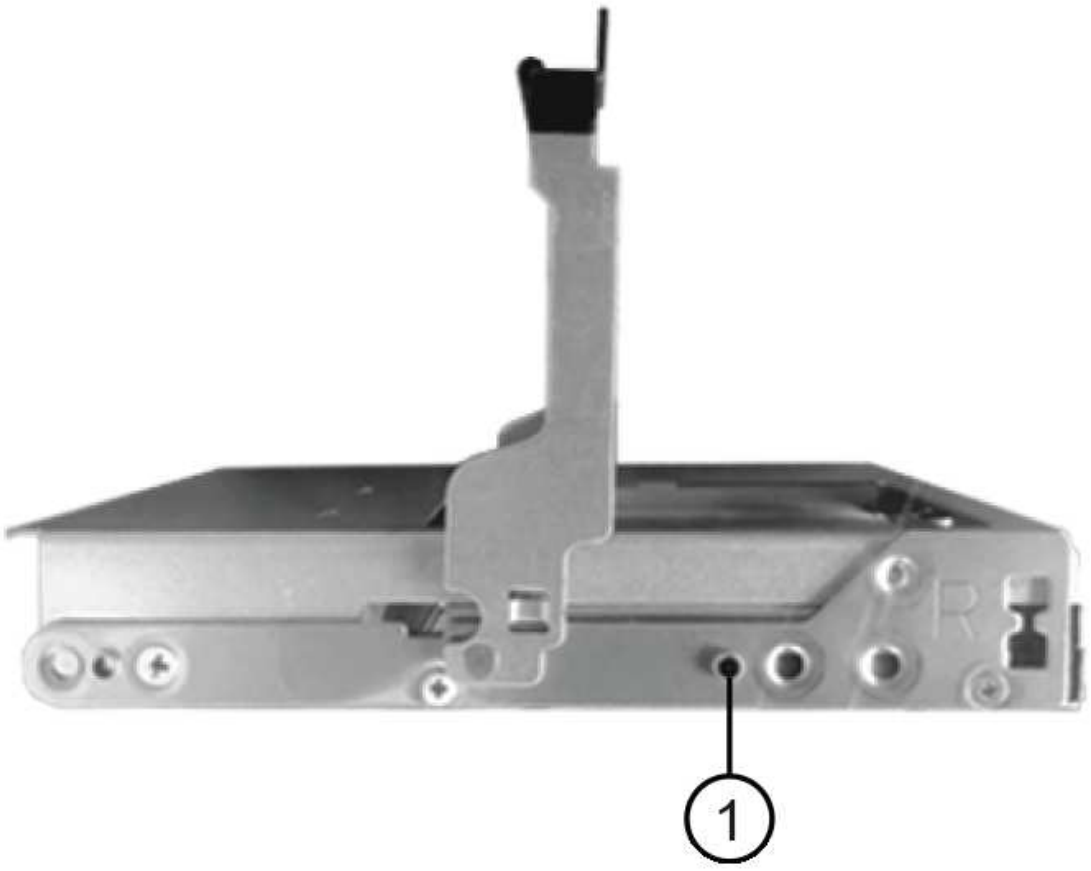

1	Raised button on the right side of the drive carrier

e. Lower the drive straight down, and then rotate the cam handle down until the drive snaps into place under the orange release latch.

f. Repeat the previous substeps for each drive in the drawer.

   You must be sure that slots 0, 3, 6, and 9 in each drawer contain drives.

g. Carefully push the drive drawer back into the enclosure.

Possible loss of data access: Never slam the drawer shut. Push the drawer in slowly to avoid jarring the drawer and causing damage to the storage array.

h. Close the drive drawer by pushing both levers towards the center.

i. Repeat these steps for each drawer in the disk shelf.

j. Attach the front bezel.

5. If you are adding multiple disk shelves, repeat this procedure for each disk shelf you are installing.

Do not power on the disk shelves at this time.

Step 2: Cable disk shelves for a new system installation

You cable disk shelf SAS connections—shelf-to-shelf (as applicable) and controller-to-shelf—to establish storage connectivity for the system.

About this task

After you cable the disk shelves, you power them on, set the shelf IDs, and complete system setup and configuration.

Before you begin

You must have met the following requirements and installed the disk shelves in the rack.

- You must have the installation and setup instructions for your platform model.

  The installation and setup instructions address the complete procedure for your system installation, setup, and configuration. You only use this procedure in conjunction with the platform installation and setup instructions if you need detailed information about installing or cabling the disk shelves to your storage system.

  Installation and setup instructions can be found by navigating to your platform in the AFF and FAS System Documentation.

- Disk shelves and controllers must not be powered on at this time.
- If you are using mini-SAS HD SAS optical cables, you must have met the rules in Mini-SAS HD SAS optical cable rules.

Steps

1. Cable the shelf-to-shelf connections within each stack if the stack has more than one disk shelf; otherwise, go to the next step:

   For a detailed explanation and examples of shelf-to-shelf “standard” cabling and shelf-to-shelf “double-wide” cabling, see shelf-to-shelf connection rules.

If…	Then…
You are cabling a multipath HA, tri-path HA, multipath, single-path HA, or single-path configuration	Cable the shelf-to-shelf connections as “standard” connectivity (using IOM ports 3 and 1): a. Beginning with the logical first shelf in the stack, connect IOM A port 3 to the next shelf’s IOM A port 1 until each IOM A in the stack is connected. b. Repeat substep a for IOM B. c. Repeat substeps a and b for each stack.
You are cabling a quad-path HA or quad-path configuration	Cable the shelf-to-shelf connections as “double-wide” connectivity: You cable the standard connectivity using IOM ports 3 and 1 and then the double-wide connectivity using IOM ports 4 and 2. a. Beginning with the logical first shelf in the stack, connect IOM A port 3 to the next shelf’s IOM A port 1 until each IOM A in the stack is connected. b. Beginning with the logical first shelf in the stack, connect IOM A port 4 to the next shelf’s IOM A port 2 until each IOM A in the stack is connected. c. Repeat substeps a and b for IOM B. d. Repeat substeps a through c for each stack.

2. Identify the controller SAS port pairs that you can use to cable the controller-to-stack connections.
   a. Check the controller-to-stack cabling worksheets and cabling examples to see whether a completed worksheet exists for your configuration.

      Controller-to-stack cabling worksheets and cabling examples for platforms with internal storage

      Controller-to-stack cabling worksheets and cabling examples for multipath HA configurations

      Controller-to-stack cabling worksheet and cabling example for a quad-path HA configuration with two quad-port SAS HBAs

   b. The next step depends on whether a completed worksheet exists for your configuration:

If…	Then…
There is a completed worksheet for your configuration	Go to the next step. You use the existing completed worksheet.

If…	Then…
There is no completed worksheet for your configuration	Fill out the appropriate controller-to-stack cabling worksheet template: Controller-to-stack cabling worksheet template for multipathed connectivity Controller-to-stack cabling worksheet template for quad-pathed connectivity

3. Cable the controller-to-stack connections using the completed worksheet.

   If needed, instructions for how to read a worksheet to cable controller-to-stack connections are available:

   How to read a worksheet to cable controller-to-stack connections for multipathed connectivity

   How to read a worksheet to cable controller-to-stack connections for quad-pathed connectivity

4. Connect the power supplies for each disk shelf:
   a. Connect the power cords first to the disk shelves, securing them in place with the power cord retainer, and then connect the power cords to different power sources for resiliency.
   b. Turn on the power supplies for each disk shelf and wait for the disk drives to spin up.
5. Set the shelf IDs and complete system setup:

   You must set shelf IDs so they are unique within the HA pair or single-controller configuration, including the internal disk shelf in applicable systems.

If…	Then…
You are manually setting shelf IDs	a. Access the shelf ID button behind the left end cap. b. Change the shelf ID to a unique ID (00 through 99). c. Power-cycle the disk shelf to make the shelf ID take effect. Wait at least 10 seconds before turning the power back on to complete the power cycle. The shelf ID blinks and the operator display panel amber LED blinks until you power cycle the disk shelf. d. Power on the controllers and complete system setup and configuration as instructed by the installation and setup instructions for your platform model.

If…	Then…
You are automatically assigning all shelf IDs in your HA pair or single-controller configuration Shelf IDs are assigned in sequential order from 00-99. For systems with an internal disk shelf, shelf ID assignment begins with the internal disk shelf.	a. Power on the controllers. b. As the controllers start booting, press `Ctrl-C` to abort the AUTOBOOT process when you see the message `Starting AUTOBOOT press Ctrl-C to abort`. If you miss the prompt and the controllers boot to ONTAP, halt both controllers, and then boot both controllers to the boot menu by entering `boot_ontap menu` at their LOADER prompt. c. Boot one controller to Maintenance mode:`boot_ontap menu` You only need to assign shelf IDs on one controller. d. From the boot menu, select option 5 for Maintenance mode. e. Automatically assign shelf IDs: `sasadmin expander_set_shelf_id -a` f. Exit Maintenance mode:`halt` g. Bring up the system by entering the following command at the LOADER prompt of both controllers:`boot_ontap` Shelf IDs appear in disk shelf digital display windows. Before you boot the system, best practice is to take this opportunity to verify cabling is correct and a root aggregate is present.. h. Complete system setup and configuration as instructed by the installation and setup instructions for your platform model.

6. If as part of system set up and configuration, you did not enable disk ownership automatic assignment, manually assign disk ownership; otherwise, go to the next step:
   a. Display all unowned disks:`storage disk show -container-type unassigned`
   b. Assign each disk:`storage disk assign -disk disk_name -owner owner_name`

      You can use the wildcard character to assign more than one disk at once.

7. Verify SAS connections are cabled correctly and there are no duplicate shelf IDs within the system by downloading and running Config Advisor as instructed by the installation and setup instructions for your platform model.

   If any SAS cabling or duplicate shelf ID errors are generated, follow the corrective actions provided.

   You can also run the `storage shelf show -fields shelf-id` command to see a list of shelf IDs already in use (and duplicates if present) in your system.

8. Verify that in-band ACP was automatically enabled. `storage shelf acp show`

   In the output, “in-band” is listed as “active” for each node.

(Optional) Step 3: Move or transport DS460C shelves

If in the future you move DS460C shelves to a different part of the data center or transport the shelves to a different location, you need to remove the drives from the drive drawers to avoid possible damage to the drive drawers and drives.

- If you saved the drive packaging materials when you installed DS460C shelves as part of your new system installation, use these to repackage the drives before moving them.

  If you did not save the packaging materials, you should place drives on cushioned surfaces or use alternate cushioned packaging. Never stack drives on top of each other.

- Before handling drives, wear an ESD wrist strap grounded to an unpainted surface on your storage enclosure chassis.

  If a wrist strap is unavailable, touch an unpainted surface on your storage enclosure chassis before handling a drive.

- You should take steps to handle drives carefully:
  - Always use two hands when removing, installing, or carrying a drive to support its weight.

    Do not place hands on the drive boards exposed on the underside of the drive carrier.

  - Be careful not to bump drives against other surfaces.
  - Drives should be kept away from magnetic devices.

    Magnetic fields can destroy all data on a drive and cause irreparable damage to the drive circuitry.

Hot-add shelves - DS212C, DS224C, or DS460C

You can hot-add one or more disk shelves with IOM12/IOM12B modules to an existing stack of disk shelves with IOM12/IOM12B modules or hot-add a stack of one or more disk shelves with IOM12/IOM12B modules directly to a SAS HBA or an onboard SAS port on the controller.

Before you begin

You must meet certain requirements and familiarize yourself with the best practices and considerations for this

procedure before hot-adding disk shelves.

- Ensure that your system meets certain requirements before hot-adding disk shelves with IOM12/IOM12B modules:
  - Your system and version of ONTAP must support the disk shelves you are hot-adding, including the IOMs, disk drives, and SAS cables. You can view what version of ONTAP you need for your shelves in the NetApp Hardware Universe.
  - Your system must have less than the maximum number of disk drives supported, by at least the number of disk shelves you plan to hot-add.

    You cannot have exceeded the maximum number of disk drives supported for your system after hot-adding disk shelves. You can view how many shelves your system can support in the NetApp Hardware Universe
  - If you are hot-adding a stack of one or more disk shelves (directly to the platform controllers), your system must have enough available PCI SAS HBA or onboard SAS ports or a combination of both.

    > If you need to install an additional PCI SAS HBA, the best practice is to use 12Gb SAS HBAs to keep controller-to-stack connectivity at 12Gbs for maximum performance.
    >
    > Using 6Gb SAS HBAs or a combination of 6Gb SAS HBAs and 12Gb SAS HBAs is supported; however, IOM12 module connections to 6Gb SAS HBAs are negotiated down to 6Gbs, resulting in lower performance.

  - Your system cannot have any SAS cabling error messages.

    Verify that your SAS connections are cabled correctly by downloading and running Config Advisor.

    You must correct any cabling errors using the corrective actions provided by the error messages.
- Familiarize yourself with the requirements and considerations for using mini-SAS HD SAS optical cables:
  - If you are using mini-SAS HD SAS optical cables or a mix of mini-SAS HD SAS optical cables and SAS copper cables in the stack of disk shelves, you must have met the rules in Mini-SAS HD SAS optical cable rules.
  - If you are hot-adding a disk shelf with mini-SAS HD SAS optical cables to a stack of disk shelves that is connected with SAS copper cables, you can temporarily have both cable types in the stack.

    After hot-adding the disk shelf, you must replace the SAS copper cables for the rest of the shelf-to-shelf connections in the stack and the controller-to-stack connections so that the stack meets the rules in Mini-SAS HD SAS optical cable rules. This means that you must have ordered the appropriate number of mini-SAS HD SAS optical cables.
- Familiarize yourself with the general considerations for completing this procedure:
  - If you are hot-adding a disk shelf with IOM12/IOM12B modules to an existing stack (of disk shelves with IOM12/IOM12B modules), you can hot-add the disk shelf to either end—the logical first or last disk shelf—of the stack.

    For single-path HA and single-path configurations, as applicable to AFF A200, AFF A220, FAS2600 series, and FAS2700 systems, you hot-add disk shelves to the end of the stack that does not have controller connections.
  - Disk shelves with IOM12/IOM12B modules must be in their own unique stack.

- This procedure assumes your configuration is using in-band ACP.

  For configurations that have in-band ACP enabled, in-band ACP is automatically enabled on hot-added disk shelves. For configurations in which in-band ACP is not enabled, hot-added disk shelves operate without any ACP functionality.

- Nondisruptive stack consolidation is not supported.

  You cannot use this procedure to hot-add disk shelves that were hot-removed from another stack in the same system when the system is powered on and serving data (I/O is in progress).

- **Best practice**: Ensure your system can recognize and utilize newly qualified disk drives by downloading the current version of the Disk Qualification Package (DQP).

  This allows you to avoid system event messages about having non-current disk drive information. You also avoid the possible prevention of disk partitioning because disk drives are not recognized. The DQP notifies you of non-current disk drive firmware.

- **Best practice**: Verify disk shelf (IOM) firmware versions, the shelf IDs already in use by your system, and obtain a snapshot of the SAS connectivity by downloading and running Config Advisor before hot-adding a disk shelf. You must also verify SAS connections are cabled correctly and that shelf IDs are unique within the HA pair or single-controller configuration by running Config Advisor after hot-adding a disk shelf.

  If any SAS cabling or duplicate shelf ID errors are generated, follow the corrective actions provided.

  You need network access to download Config Advisor.

- **Best practice**: Ensure your system has the current versions of disk shelf (IOM) firmware and disk drive firmware before adding new disk shelves, shelf FRU components, or SAS cables. You can visit the NetApp Support Site to download disk shelf firmware and download disk drive firmware.
- Familiarize yourself with the considerations for properly handling SAS cables:
  - Visually inspect the SAS port to verify the proper orientation of the connector before plugging it in.

    The SAS cable connectors are keyed. When oriented correctly into a SAS port, the connector clicks into place and if the disk shelf power is on at the time, the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector).

    For controllers, the orientation of SAS ports can vary depending on the platform model; therefore, the correct orientation of the SAS cable connector varies.

  - To prevent degraded performance, do not twist, fold, pinch, or step on the cables.

    Cables have a minimum bend radius. Cable manufacturer specifications define the minimum bend radius; however, a general guideline for minimum bend radius is 10 times the cable diameter.

  - Using Velcro wraps instead of tie-wraps to bundle and secure system cables allows for easier cable adjustments.
- Familiarize yourself with the considerations for properly handling DS460C drives:
  - The drives are packaged separately from the shelf chassis.

    You should take inventory of the drives.

- After you unpack the drives, you should save the packaging materials for future use.

  Possible loss of data access: If in the future, you move the shelf to a different part of the data center or transport the shelf to a different location, you need to remove the drives from the drive drawers to avoid possible damage to the drive drawers and drives.

  Keep disk drives in their ESD bag until you are ready to install them.

- When handling the drives, always wear an ESD wrist strap grounded to an unpainted surface on your storage enclosure chassis to prevent static discharges.

  If a wrist strap is unavailable, touch an unpainted surface on your storage enclosure chassis before handling the disk drive.

Step 1: Install disk shelves for a hot-add

For each disk shelf you are hot-adding, you install the disk shelf into a rack, connect the power cords, power on the disk shelf, and set the disk shelf ID before cabling the SAS connections.

Steps

1. Install the rack mount kit (for two-post or four-post rack installations) that came with your disk shelf using the installation flyer that came with the kit.

   If you are installing multiple disk shelves, you should install them from the bottom to the top of the rack for the best stability.

   Do not flange-mount the disk shelf into a telco-type rack; the disk shelf's weight can cause it to collapse in the rack under its own weight.

2. Install and secure the disk shelf onto the support brackets and rack using the installation flyer that came with the kit.

   To make a disk shelf lighter and easier to maneuver, remove the power supplies and I/O modules (IOMs).

   Although the drives in DS460C shelves are packaged separately, which makes the shelf lighter, an empty DS460C shelf still weighs approximately 132 lb (60kg). It is recommended that you use a mechanized lift or four people using the lift handles to safely move an empty DS460C shelf.

   Your DS460C shipment includes four detachable lift handles (two for each side). To use the lift handles, you install them by inserting the tabs of the handles into the slots in the side of the shelf and pushing up until they click into place. Then, as you slide the disk shelf onto the rails, you detach one set of handles at a time using the thumb latch. The following illustration shows how to attach a lift handle.

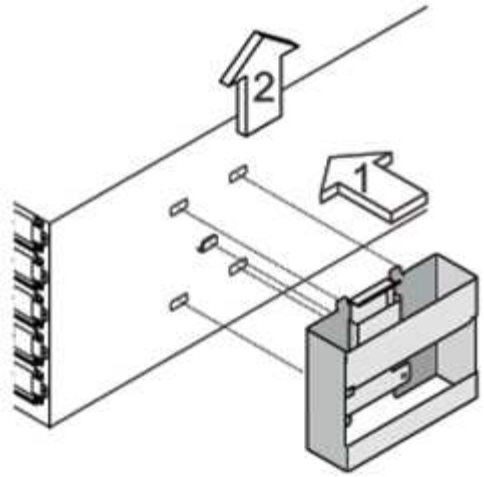

3. Reinstall any power supplies and IOMs you removed prior to installing your disk shelf into the rack.
4. If you are installing a DS460C disk shelf, install the drives into the drive drawers. Otherwise, go to the next step.

> Always wear an ESD wrist strap grounded to an unpainted surface on your storage enclosure chassis to prevent static discharges.
>
> If a wrist strap is unavailable, touch an unpainted surface on your storage enclosure chassis before handling the disk drive.

If you purchased a partially populated shelf, meaning that the shelf has less than the 60 drives it supports, install the drives as follows in each drawer:

- Install the first four drives into the front slots (0, 3, 6, and 9).

> **Risk of equipment malfunction:** To allow for proper air flow and prevent overheating, always install the first four drives into the front slots (0, 3, 6, and 9).

- For the remaining drives, evenly distribute them across each drawer.

  The following illustration shows how the drives are numbered from 0 to 11 in each drive drawer within the shelf.

a. Open the top drawer of the shelf.
b. Remove a drive from its ESD bag.
c. Raise the cam handle on the drive to vertical.
d. Align the two raised buttons on each side of the drive carrier with the matching gap in the drive channel on the drive drawer.

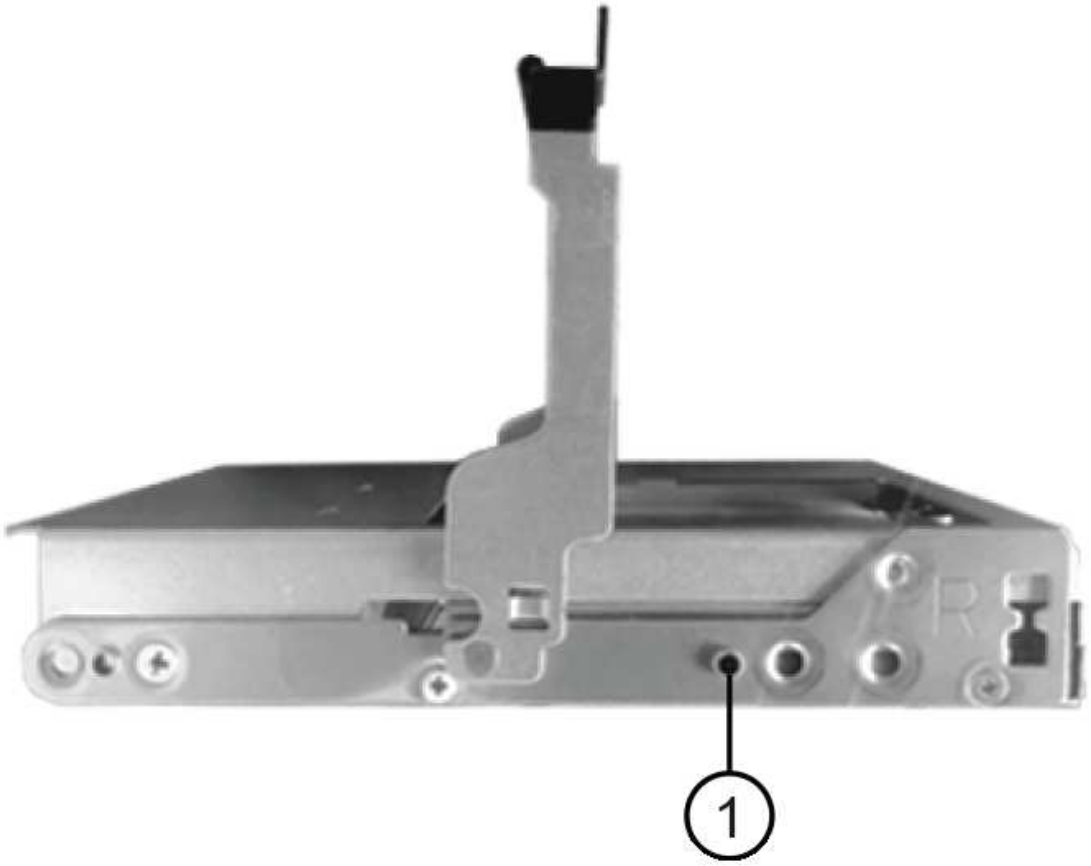

1	Raised button on the right side of the drive carrier

e. Lower the drive straight down, and then rotate the cam handle down until the drive snaps into place under the orange release latch.

f. Repeat the previous substeps for each drive in the drawer.

You must be sure that slots 0, 3, 6, and 9 in each drawer contain drives.

g. Carefully push the drive drawer back into the enclosure.
+s

Possible loss of data access: Never slam the drawer shut. Push the drawer in slowly to avoid jarring the drawer and causing damage to the storage array.

h. Close the drive drawer by pushing both levers towards the center.

i. Repeat these steps for each drawer in the disk shelf.

j. Attach the front bezel.

5. If you are adding multiple disk shelves, repeat the previous steps for each disk shelf you are installing.
6. Connect the power supplies for each disk shelf:
   a. Connect the power cords first to the disk shelves, securing them in place with the power cord retainer, and then connect the power cords to different power sources for resiliency.
   b. Turn on the power supplies for each disk shelf and wait for the disk drives to spin up.
7. Set the shelf ID for each disk shelf you are hot-adding to an ID that is unique within the HA pair or single-controller configuration.

   If you have a platform model with an internal disk shelf, shelf IDs must be unique across the internal disk shelf and externally attached disk shelves.

   You can use the following substeps to change shelf IDs. For more detailed instructions, use Change a shelf ID.

   a. If needed, verify shelf IDs already in use by running Config Advisor.

   You can also run the `storage shelf show -fields shelf-id` command to see a list of shelf IDs already in use (and duplicates if present) in your system.

   b. Access the shelf ID button behind the left end cap.

c. Change the shelf ID to a valid ID (00 through 99).

d. Power-cycle the disk shelf to make the shelf ID take effect.

   Wait at least 10 seconds before turning the power back on to complete the power cycle.

   The shelf ID blinks and the operator display panel amber LED blinks until you power cycle the disk shelf.

e. Repeat substeps a through d for each disk shelf you are hot-adding.

Step 2: Cable disk shelves for a hot-add

You cable the SAS connections (shelf-to-shelf and controller-to-stack) as applicable for hot-added disk shelves so they have connectivity to the system.

About this task

- For an explanation and examples of shelf-to-shelf “standard” cabling and shelf-to-shelf “double-wide” cabling, see Shelf-to-shelf SAS connection rules.
- For instructions about how to read a worksheet to cable controller-to-stack connections, see How to read a worksheet to cable controller-to-stack connections for multipathed connectivity or How to read a worksheet to cable controller-to-stack connections for quad-pathed connectivity.
- After you have cabled the hot-added disk shelves, ONTAP recognizes them: disk ownership is assigned if disk ownership automatic assignment is enabled; disk shelf (IOM) firmware and disk drive firmware should automatically update if needed; and if in-band ACP is enabled on your configuration, it is automatically enabled on the hot-added disk shelves.

  Firmware updates can take up to 30 minutes.

Before you begin

- You must have met the requirements for completing this procedure and installed, powered on, and set shelf IDs for each disk shelf as instructed in Install disk shelves with IOM12 modules for a hot-add.

Steps

1. If you want to manually assign disk ownership for the disk shelves you are hot-adding, you need to disable disk ownership automatic assignment if it is enabled; otherwise, go to the next step.

   You need to manually assign disk ownership if disks in the stack are owned by both controllers in an HA pair.

   You disable disk ownership automatic assignment before cabling the hot-added disk shelves and then later, in step 7, you reenable it after cabling the hot-added disk shelves.

   a. Verify if disk ownership automatic assignment is enabled:`storage disk option show`

      If you have an HA pair, you can enter the command at the console of either controller.

      If disk ownership automatic assignment is enabled, the output shows “on” (for each controller) in the “Auto Assign” column.

   b. If disk ownership automatic assignment is enabled, you need to disable it:`storage disk option modify -node _node_nam_e -autoassign off`

You need to disable disk ownership automatic assignment on both controllers in an HA pair.

2. If you are hot-adding a stack of disk shelves directly to a controller, complete the following substeps; otherwise, go to step 3.

   a. If the stack you are hot-adding has more than one disk shelf, cable the shelf-to-shelf connections; otherwise, go to substep b.

If…	Then…
You are cabling a stack with multipath HA, tri-path HA, multipath, single-path HA, or single-path connectivity to the controllers	Cable the shelf-to-shelf connections as “standard” connectivity (using IOM ports 3 and 1): i. Beginning with the logical first shelf in the stack, connect IOM A port 3 to the next shelf’s IOM A port 1 until each IOM A in the stack is connected. ii. Repeat substep i for IOM B.
You are cabling a stack with quad-path HA or quad-path connectivity to the controllers	Cable the shelf-to-shelf connections as “double-wide” connectivity: You cable the standard connectivity using IOM ports 3 and 1 and then the double-wide connectivity using IOM ports 4 and 2. i. Beginning with the logical first shelf in the stack, connect IOM A port 3 to the next shelf’s IOM A port 1 until each IOM A in the stack is connected. ii. Beginning with the logical first shelf in the stack, connect IOM A port 4 to the next shelf’s IOM A port 2 until each IOM A in the stack is connected. iii. Repeat substeps i and ii for IOM B.

   b. Check the controller-to-stack cabling worksheets and cabling examples to see whether a completed worksheet exists for your configuration.

      Controller-to-stack cabling worksheets and cabling examples for platforms with internal storage

      Controller-to-stack cabling worksheets and cabling examples for multipath HA configurations

      Controller-to-stack cabling worksheet and cabling example for a quad-path HA configuration with two quad-port SAS HBAs

   c. If there is a completed worksheet for your configuration, cable the controller-to-stack connections using the completed worksheet; otherwise, go to the next substep.

   d. If there is no completed worksheet for your configuration, fill out the appropriate worksheet template, and then cable the controller-to-stack connections using the completed worksheet.

      Controller-to-stack cabling worksheet template for multipathed connectivity

      Controller-to-stack cabling worksheet template for quad-pathed connectivity

e. Verify that all cables are securely fastened.

3. If you are hot-adding one or more disk shelves to an end (the logical first or last disk shelf) of an existing stack, complete the applicable substeps for your configuration; otherwise, go to the next step.

> Make sure that you wait at least 70 seconds between disconnecting a cable and reconnecting it, and if you are replacing a cable for a longer one.

If you are…	Then…
Hot-adding a disk shelf to an end of a stack that has multipath HA, tri-path HA, multipath, quad-path HA, or quad-path connectivity to the controllers	a. Disconnect any cables from IOM A of the disk shelf at the end of the stack that are connected to any controllers; otherwise, go to substep e. Leave the other end of these cables connected to the controllers, or replace cables with longer cables if needed. b. Cable the shelf-to-shelf connection(s) between IOM A of the disk shelf at the end of the stack and IOM A of the disk shelf you are hot-adding. c. Reconnect any cables that you removed in substep a to the same port(s) on IOM A of the disk shelf you are hot-adding; otherwise, go to the next substep. d. Verify that all cables are securely fastened. e. Repeat substeps a through d for IOM B; otherwise, go to Step 4.
Hot-adding a disk shelf to an end of the stack in a single-path HA or single-path configuration, as applicable to AFF A200, AFF A220, FAS2600 series and FAS2700 systems. These instructions are for hot-adding to the end of the stack that does not have controller-to-stack connections.	a. Cable the shelf-to-shelf connection between IOM A of the disk shelf in the stack and IOM A of the disk shelf you are hot-adding. b. Verify that the cable is securely fastened. c. Repeat applicable substeps for IOM B.

4. If you hot-added a disk shelf with mini-SAS HD SAS optical cables to a stack of disk shelves connected with SAS copper cables, replace the SAS copper cables; otherwise, go to the next step.

   Replace cables one at a time and make sure that you wait at least 70 seconds between disconnecting a cable and connecting a new one.

5. Verify that your SAS connections are cabled correctly by downloading and running Config Advisor.

   If any SAS cabling errors are generated, follow the corrective actions provided.

6. Verify SAS connectivity for each hot-added disk shelf: `storage shelf show -shelf` *`shelf_name`* `-connectivity`

   You must run this command for each disk shelf you hot-added.

For example, the following output shows hot-added disk shelf 2.5 is connected to initiator ports 1a and 0d (port pair 1a/0d) on each controller (in a FAS8080 multipath HA configuration with one quad-port SAS HBA):

```
cluster1::> storage shelf show -shelf 2.5 -connectivity

           Shelf Name: 2.5
             Stack ID: 2
             Shelf ID: 5
            Shelf UID: 40:0a:09:70:02:2a:2b
        Serial Number: 101033373
          Module Type: IOM12
                Model: DS224C
         Shelf Vendor: NETAPP
           Disk Count: 24
      Connection Type: SAS
          Shelf State: Online
               Status: Normal

Paths:

Controller     Initiator   Initiator Side Switch Port   Target Side
Switch Port   Target Port   TPGN
------------   ---------   --------------------------
-----------------------   -----------   ------
stor-8080-1    1a           -                           -
-             -
stor-8080-1    0d           -                           -
-             -
stor-8080-2    1a           -                           -
-             -
stor-8080-2    0d           -                           -
-             -

Errors:
------
-
```

7. If you disabled disk ownership automatic assignment in Step 1, manually assign disk ownership, and then reenable disk ownership automatic assignment if needed:
   a. Display all unowned disks:`storage disk show -container-type unassigned`
   b. Assign each disk:`storage disk assign -disk` *`disk_name`* `-owner` *`owner_name`*

      You can use the wildcard character to assign more than one disk at once.

   c. Reenable disk ownership automatic assignment if needed:`storage disk option modify -node`

node_name -autoassign on

You need to reenable disk ownership automatic assignment on both controllers in an HA pair.

8. If your configuration is running in-band ACP, verify that in-band ACP was automatically enabled on hot-added disk shelves: `storage shelf acp show`

   In the output, “in-band” is listed as “active” for each node.

(Optional) Step 3: Move or transport DS460C shelves

If in the future you move DS460C shelves to a different part of the data center or transport the shelves to a different location, you need to remove the drives from the drive drawers to avoid possible damage to the drive drawers and drives.

- If you saved the drive packaging materials when you installed DS460C shelves as part of your shelf hot-add, use these to repackage the drives before moving them.

  If you did not save the packaging materials, you should place drives on cushioned surfaces or use alternate cushioned packaging. Never stack drives on top of each other.

- Before handling drives, wear an ESD wrist strap grounded to an unpainted surface on your storage enclosure chassis.

  If a wrist strap is unavailable, touch an unpainted surface on your storage enclosure chassis before handling a drive.

- You should take steps to handle drives carefully:
  - Always use two hands when removing, installing, or carrying a drive to support its weight.

    Do not place hands on the drive boards exposed on the underside of the drive carrier.

  - Be careful not to bump drives against other surfaces.
  - Drives should be kept away from magnetic devices.

    Magnetic fields can destroy all data on a drive and cause irreparable damage to the drive circuitry.

Change a shelf ID - DS212C, DS224C, or DS460C

You can change a shelf ID in a system with IOM12/IOM12B modules when ONTAP is not yet running or when hot-adding a shelf prior to it being cabled to the system. You can also change a shelf ID when ONTAP is up and running (controller modules are available to serve data) and all drives in the shelf are unowned, spares, or part of offlined aggregate(s).

About this task

- A valid shelf ID is 00 through 99.
- Shelf IDs must be unique within an HA pair or single-controller configuration.

If you have a platform with internal storage, shelf IDs must be unique across the internal disk shelf and any externally attached disk shelves.

- You must power cycle a shelf in order for the shelf ID to take effect.

  The amount of time you wait before turning the power back on depends on the state of ONTAP, as described later in this procedure.

Before you begin

- If ONTAP is up and running (controller modules are available to serve data), you must have verified that all drives in the shelf are unowned, spares, or part of offlined aggregate(s).

  You can verify the state of the drives by using the `storage disk show -shelf shelf_number` command. Output in the Container Type column should display spare or broken if it is a failed drive. Additionally, the Container Name and Owner columns should have a dash.

- You can verify shelf IDs already in use in your system by running Active IQ Config Advisor or using the `storage shelf show -fields shelf-id` command. You can download and access the Active IQ Config Advisor on the NetApp Support Site.

Steps

1. Turn on the power to the disk shelf if it is not already on.
2. Remove the left end cap to locate the button near the shelf LEDs.
3. Change the first number of the shelf ID by pressing and holding the orange button until the first number on the digital display blinks, which can take up to three seconds.

   If the ID takes longer than three seconds to blink, press the button again, making sure to press it in all the way.

   This activates the disk shelf ID programming mode.

4. Press the button to advance the number until you reach the desired number from 0 to 9.

   The first number continues to blink.

5. Change the second number of the shelf ID by pressing and holding the button until the second number on the digital display blinks, which can take up to three seconds.

   The first number on the digital display stops blinking.

6. Press the button to advance the number until you reach the desired number from 1 to 9.

   The second number continues to blink.

7. Lock in the desired number and exit the programming mode by pressing and holding the button until the second number stops blinking, which can take up to three seconds.

   Both numbers on the digital display start blinking and the amber LED on the operator display panel illuminates after about five seconds, alerting you that the pending disk shelf ID has not yet taken effect.

8. Power cycle the disk shelf to make the shelf ID take effect.

   You must turn off both power switches, wait the appropriate amount of time, and then turn them back on to

complete the power cycle.

- If ONTAP is not yet running or you are hot-adding a shelf (that has not yet been cabled to the system), wait at least 10 seconds.
- If ONTAP is running (controllers are available to serve data), and all disk drives in the shelf are unowned, spares, or part of offlined aggregate(s), wait at least 120 seconds.

  This time allows ONTAP to properly delete the old shelf address and update the copy of the new shelf address.

9. Replace the left end cap.
10. Repeat steps 1 through 9 for each additional disk shelf.
11. Verify that your system does not have duplicate shelf IDs.

    When two or more disk shelves have the same ID, the system assigns the duplicate disk shelf a soft ID number equal to or greater than 100. You must change the soft ID (duplicate) number.

    a. Run Active IQ Config Advisor to check for duplicate shelf ID alerts or run the `storage shelf show -fields shelf-id` command to see a list of shelf IDs already in use including any duplicate IDs.

    b. If your system has any duplicate shelf IDs, change the duplicate shelf IDs by repeating this procedure.

SAS cabling rules, worksheets, and examples

Overview of cabling rules - DS212C, DS224C, or DS460C

To help you cable your SAS drive shelves with IOM12/IOM12B modules to your storage system, you can use any of the available SAS cabling rules, worksheets, and examples content as needed.

SAS cabling rules and concepts

- Configurations
- Controller slot numbering
- Shelf-to-shelf connections
- Controller-to-stack connections
- Mini-SAS HD SAS optical cables
- Tri-path HA connectivity

Cabling worksheets and examples

- Multipath HA configurations
- Platforms with internal storage
- Quad-path HA configurations

Cabling worksheet templates

- Multipathed connectivity
- Quad-pathed connectivity

- How to read a worksheet for multipathed connectivity
- How to read a worksheet for quad-pathed connectivity

SAS cabling rules and concepts - DS212C, DS224C, or DS460C

Disk shelves with IOM12/IOM12B modules can be cabled in HA pair and single-controller configurations (for supported platforms) by applying the SAS cabling rules: configuration rules, controller slot numbering rules, shelf-to-shelf connection rules, controller-to-stack connection rules, and if applicable, mini-SAS HD SAS optical cable rules.

> The SAS cabling rules regarding controller slot numbering rules, shelf-to-shelf connection rules, and controller-to-stack connection rules described in this guide are the same rules that apply to all SAS disk shelves, whether they have IOM12 or IOM12B modules. However, the information in this guide is specific to the unique characteristics of disk shelves with IOM12/IOM12B modules and their use in supported configurations.

The SAS cabling rules regarding configuration rules and mini-SAS HD SAS optical cable rules described in this guide are specific to disk shelves with IOM12/IOM12B modules.

The SAS cabling rules described in this guide balance SAS cabling between the on-board SAS ports and host bus adapter SAS ports to provide highly available storage controller configurations and meet the following goals:

- Provide a single, easily understood universal algorithm for all SAS products and configurations
- Yield the same physical cabling when generating the Bill of Materials (BOM), followed in the factory, and in the field
- Are verifiable by configuration-checking software and tools
- Provide maximum possible resilience to maintain availability and minimize the reliance on controller takeovers

You should avoid deviating from the rules; deviations might reduce reliability, universality, and commonality.

Configuration rules

Disk shelves with IOM12/IOM12B modules are supported on specific types of HA pair and single-controller configurations.

> For current information about supported cabling configurations for your platform model, see the Hardware Universe.
>
> NetApp Hardware Universe

- HA pair configurations must be cabled as multipath HA or quad-path HA configurations with the following exceptions:
  - Platforms with internal storage do not support quad-path HA connectivity.
  - A FAS2820 HA pair can be cabled as tri-path HA.

    Information about the FAS2820 connectivity can be found in the Tri-path HA connectivity section.
  - Platforms with internal storage can be cabled as single-path HA configurations (from port 0b/0b1 to

external shelves) to support connectivity to an external SAS tape backup device (from port 0a).

> For FAS2820 HA pairs, although the cabling to external shelves is single-path HA, because of each controller's internal connection of port 0b to its local expander (IOM12G) and port 0c to its partner's expander, the HA pair configuration is multipath HA.

- Single-controller configurations must be cabled as multipath or quad-path configurations, with the following exceptions:
  - FAS2600 series single-controller configurations can be cabled as single-path configurations.

    Because the internal storage uses single-path connectivity, ONTAP issues occasional warnings that mixed paths are detected. To avoid these warnings, you can use single-path connectivity to the external disk shelves. Additionally, you can use single-path connectivity when an external SAS tape backup device is used.

  - FAS2600 series single-controller configurations do not support quad-path connectivity.

Controller slot numbering rules

For the purpose of applying cabling rules across all supported HA pairs and single-controller configurations, a controller slot numbering convention is used.

- For all HA pairs and single-controller configurations, the following applies:
  - A SAS HBA in a physical PCI slot is defined as occupying PCI slot 1, 2, 3, and so on regardless of the slot's physical label on a controller.

    For example, if SAS HBAs occupied physical PCI slots 3, 5, and 7, they would be designated as slots 1, 2, and 3 for the purpose of applying the SAS cabling rules.

  - An onboard SAS HBA is defined as occupying PCI slot 0 just as it is labeled on a controller.
  - Each port in each slot is defined just as it is labeled on a controller.
    For example, slot 0 with two ports is referred to as 0a and 0b. Slot 1 with four ports is referred to as 1a, 1b, 1c, and 1d.

    In this document, slots and the slot ports are depicted as follows:

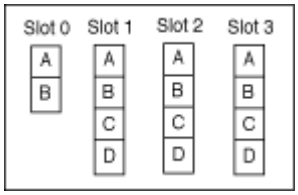

Shelf-to-shelf connection rules

When you have more than one disk shelf in a stack of disk shelves, they connect to each other through each SAS domain (IOM A and IOM B) using the applicable "standard" or "double-wide" shelf-to-shelf cabling. Your use of "standard" or "double-wide" shelf-to-shelf cabling depends on the configuration you have.

Standard shelf-to-shelf connectivity

- Standard shelf-to-shelf connectivity is used in any stack of disk shelves with more than one disk shelf.

One cable connection is needed between disk shelves in each domain—domain A (IOM A) and domain B (IOM B).

- Best practice is to use IOM ports 3 and 1 for standard shelf-to-shelf connectivity.

  From the logical first shelf to the logical last shelf in a stack, you connect IOM port 3 to the next shelf's IOM port 1 in domain A and then domain B.

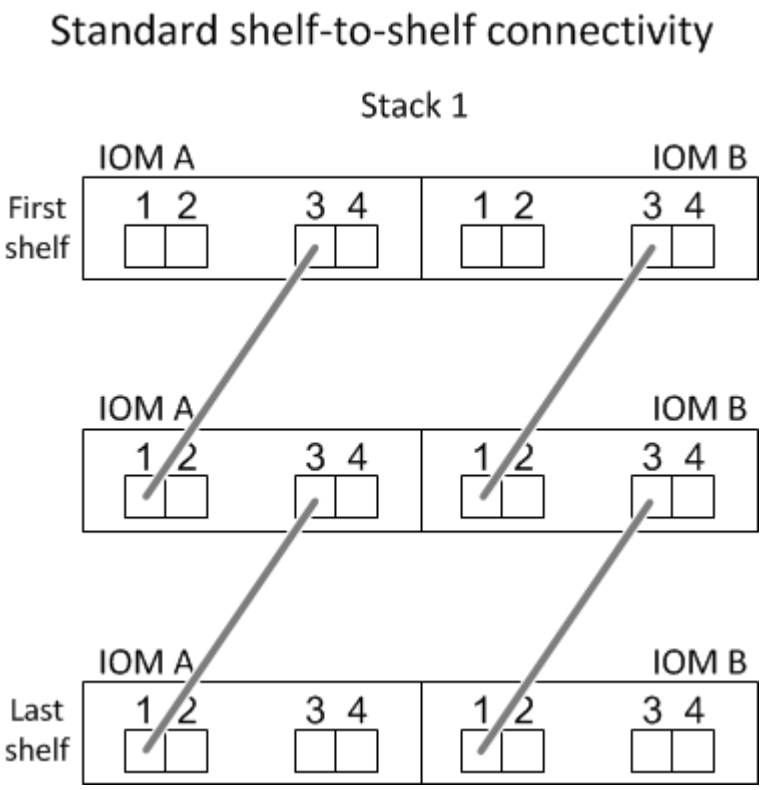

Double-wide shelf-to-shelf connectivity

- Double-wide shelf-to-shelf connectivity is used in quad-pathed (quad-path HA and quad-path) configurations.
- Double-wide shelf-to-shelf connectivity requires two cable connections between disk shelves in each domain—domain A (IOM A) and domain B (IOM B).

  The first cable connection is cabled as standard shelf-to-shelf connectivity (using IOM ports 3 and 1); the second cable connection is cabled as double-wide shelf-to-shelf connectivity (using IOM ports 4 and 2).

  From the logical first shelf to the logical last shelf in a stack, you connect IOM port 3 to the next shelf's IOM port 1 in domain A and then domain B. From the logical first shelf to the logical last shelf in a stack, you connect IOM port 4 to the next shelf's IOM port 2 in domain A and then domain B. (IOM ports cabled as double-wide connectivity are shown with blue.)

Double-wide shelf-to-shelf connectivity

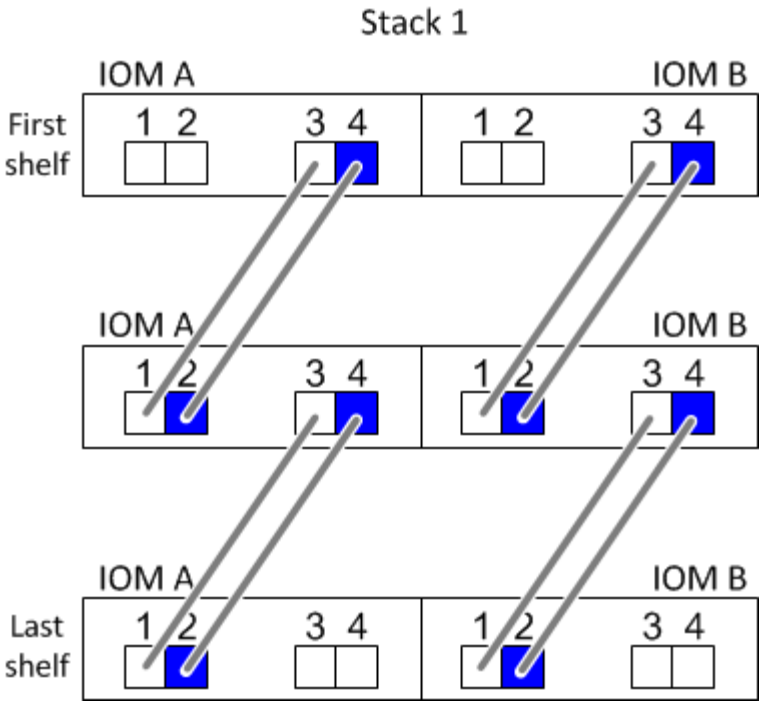

Controller-to-stack connection rules

You can correctly cable the SAS connections from each controller to each stack in an HA pair or in a single-controller configuration by understanding that SAS disk shelves use software-based disk ownership, how controller ports A/C and B/D are connected to stacks, how controller ports A/C and B/D are organized into port pairs, and how platforms with internal storage have their controller ports connected to stacks.

SAS disk shelf software-based disk ownership rule

SAS disk shelves use software-based disk ownership (not hardware-based disk ownership). This means that disk drive ownership is stored on the disk drive rather than it being determined by the topology of the storage system's physical connections (as it is for hardware-based disk ownership). Specifically, disk drive ownership is assigned by ONTAP (automatically or by CLI commands), not by how you cable the controller-to-stack connections.

SAS disk shelves should never be cabled using the hardware-based disk ownership scheme.

Controller A and C port connection rules (for platforms without internal storage)

- A and C ports are always the primary paths to a stack.
- A and C ports always connect to the logical first disk shelf in a stack.
- A and C ports always connect to disk shelf IOM ports 1 and 2.

  IOM port 2 is only used for quad-path HA and quad-path configurations.
- Controller 1 A and C ports always connect to IOM A (domain A).
- Controller 2 A and C ports always connect to IOM B (domain B).

The following illustration highlights how controller ports A and C connect in a multipath HA configuration with

one quad-port HBA and two stacks of disk shelves. Connections to stack 1 are shown in blue. Connections to stack 2 are shown in orange.

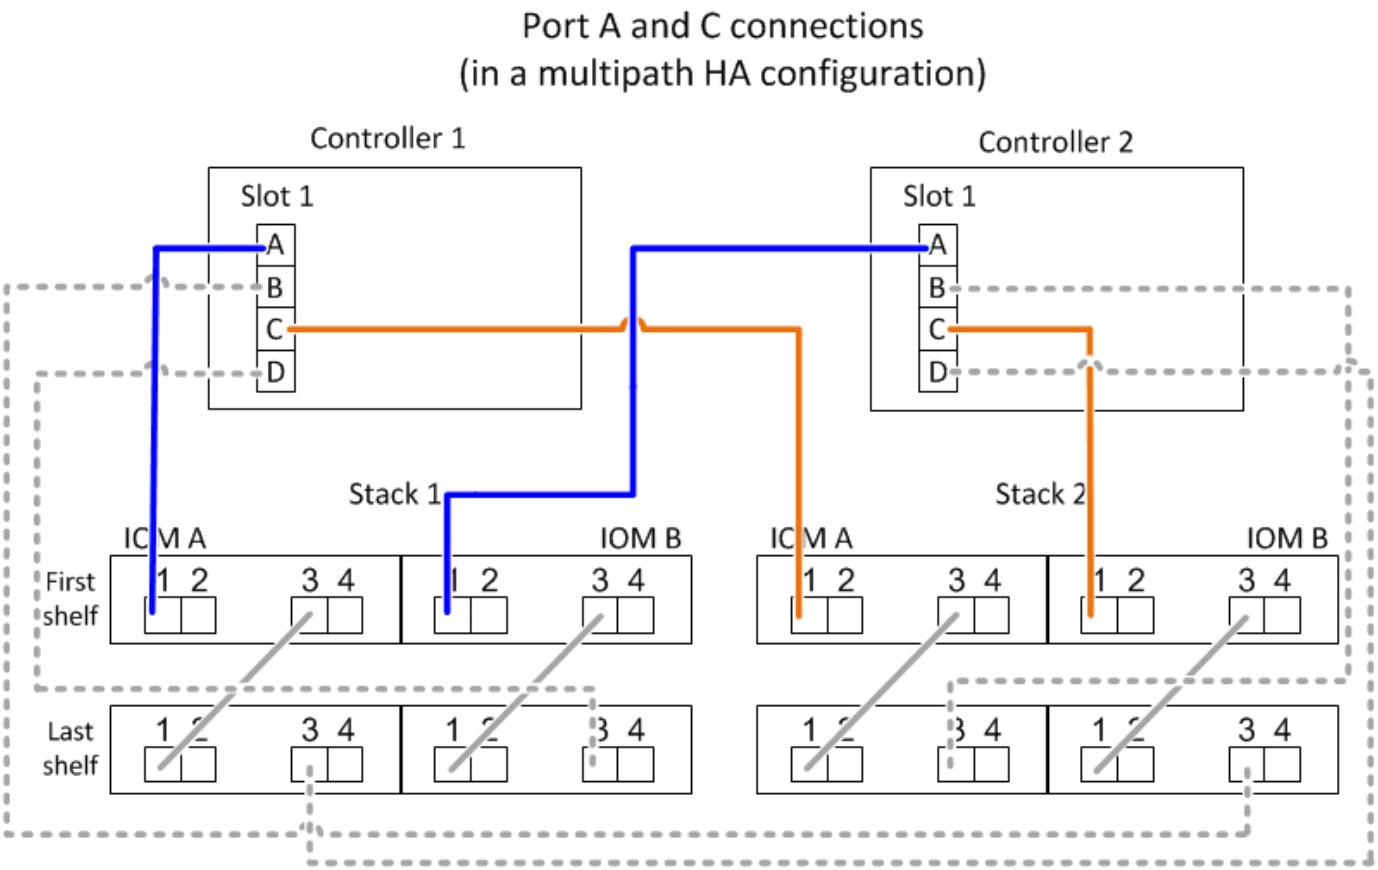

Controller B and D port connection rules (for platforms without internal storage)

- B and D ports are always the secondary paths to a stack.
- B and D ports always connect to the logical last disk shelf in a stack.
- B and D ports always connect to disk shelf IOM ports 3 and 4.

  IOM port 4 is only used for quad-path HA and quad-path configurations.

- Controller 1 B and D ports always connect to IOM B (domain B).
- Controller 2 B and D ports always connect to IOM A (domain A).
- B and D ports are connected to the stacks by offsetting the order of the PCI slots by one so that the first port on the first slot is cabled last.

The following illustration highlights how controller ports B and D connect in a multipath HA configuration with one quad-port HBA and two stacks of disk shelves. Connections to stack 1 are shown in blue. Connections to stack 2 are shown in orange.

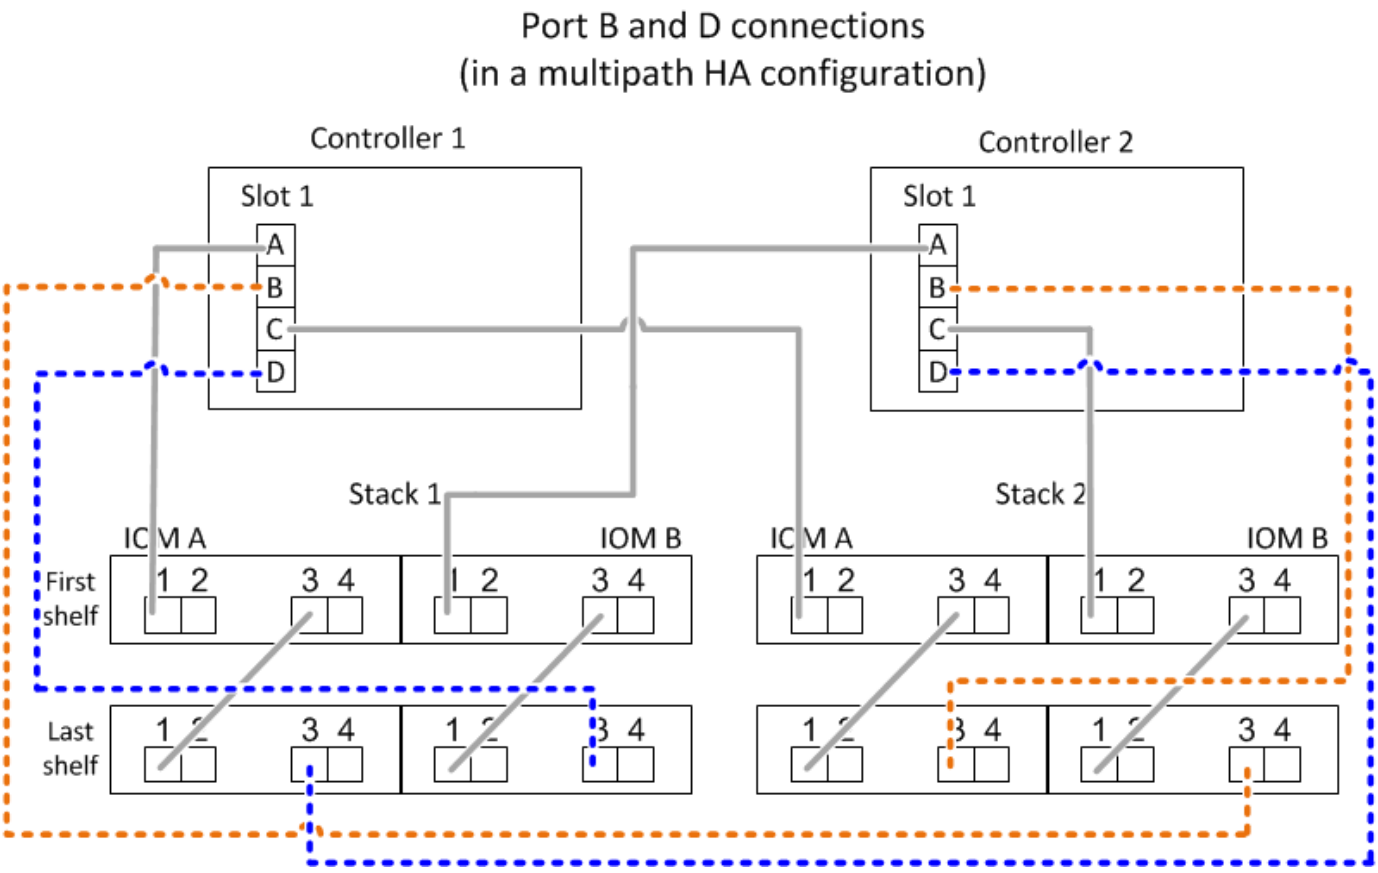

Port pair connection rules (for platforms without internal storage)

Controller SAS ports A, B, C, and D are organized into port pairs using a method that leverages all of the SAS ports for system resiliency and consistency when cabling controller-to-stack connections in HA pair and single-controller configurations.

- Port pairs consist of a controller A or C SAS port and a controller B or D SAS port.

  A and C SAS ports connect to the logical first shelf in a stack. B and D SAS ports connect to the logical last shelf in a stack.

- Port pairs use all SAS ports on each controller in your system.

  You increase system resiliency by incorporating all SAS ports (on an HBA in a physical PCI slot [slot 1-N] and on board the controller [slot 0]) into port pairs. Do not exclude any SAS ports.

- Port pairs are identified and organized as follows:
  a. List A ports and then C ports in sequence of slots (0,1, 2, 3, and so on).

     For example: 1a, 2a, 3a, 1c, 2c, 3c

  b. List B ports and then D ports in sequence of slots (0,1, 2, 3, and so on).

     For example: 1b, 2b, 3b, 1d, 2d, 3d

  c. Rewrite the D and B port list so that the first port in the list is moved to the end of the list.

For example: 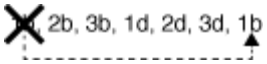

Offsetting the order of the slots by one balances port pairs across multiple slots (physical PCI slots and on board slots) when more than one slot of SAS ports is available; therefore, preventing a stack from being cabled to a single SAS HBA.

d. Pair the A and C ports (listed in step 1) to the D and B ports (listed in step 2) in the order that they are listed.

For example: 1a/2b, 2a/3b, 3a/1d,1c/2d, 2c/3d, 3c/1b.

> For an HA pair, the list of port pairs you identify for the first controller is also applicable to the second controller.

- When cabling your system, you can use port pairs in the order in which you identified them or you can skip port pairs:
  - Use port pairs in the order in which you identified (listed) them when all port pairs are needed to cable the stacks in your system.

    For example, if you identified six port pairs for your system and you have six stacks to cable as multipath, you cable the port pairs in the order in which you listed them:

    1a/2b, 2a/3b, 3a/1d,1c/2d, 2c/3d, 3c/1b

  - Skip port pairs (use every other port pair) when not all port pairs are needed to cable the stacks in your system.

    For example, if you identified six port pairs for your system and you have three stacks to cable as multipath, you cable every other port pair in your list:

    1a/2b, ~~2a/3b~~, 3a/1d, ~~1c/2d~~, 2c/3d, ~~3c/1b~~

> When you have more port pairs than you need to cable the stacks in your system, the best practice is to skip port pairs to optimize the SAS ports on your system. By optimizing SAS ports, you optimize your system's performance.

Controller-to-stack cabling worksheets are convenient tools for identifying and organizing port pairs so that you can cable the controller-to-stack connections for your HA pair or single-controller configuration.

Controller-to-stack cabling worksheet template for multipathed connectivity

Controller-to-stack cabling worksheet template for quad-pathed connectivity

Controller 0b/0b1 and 0a port connection rules for platforms with internal storage

Platforms with internal storage have a unique set of connection rules because each controller must maintain same domain connectivity between the internal storage (port 0b/0b1) and the stack. This means that when a controller is located in slot A of the chassis (controller 1) it is in domain A (IOM A) and therefore port 0b/0b1 must connect to IOM A in the stack. When a controller is located in slot B of the chassis (controller 2) it is in domain B (IOM B) and therefore port 0b/0b1 must connect to IOM B in the stack.

FAS25XX platforms are not addressed in this content.

If you do not connect the 0b/0b1 port to the correct domain (cross-connect domains), you expose your system to resiliency issues that prevent you from performing nondisruptive procedures safely.

- Controller 0b/0b1 port (internal storage port):
  - Controller 1 0b/0b1 port always connects to IOM A (domain A).
  - Controller 2 0b/0b1 port always connects to IOM B (domain B).
  - Port 0b/0b1 is always the primary path.
  - Port 0b/0b1 always connects to the logical last disk shelf in a stack.
  - Port 0b/0b1 always connect to disk shelf IOM port 3.
- Controller 0a port (internal HBA port):
  - Controller 1 0a port always connects to IOM B (domain B).
  - Controller 2 0a port always connects to IOM A (domain A).
  - Port 0a is always the secondary path.
  - Port 0a always connects to the logical first disk shelf in a stack.
  - Port 0a always connect to disk shelf IOM port 1.

The following illustration highlights internal storage port (0b/0b1) domain connectivity to an external stack of shelves:

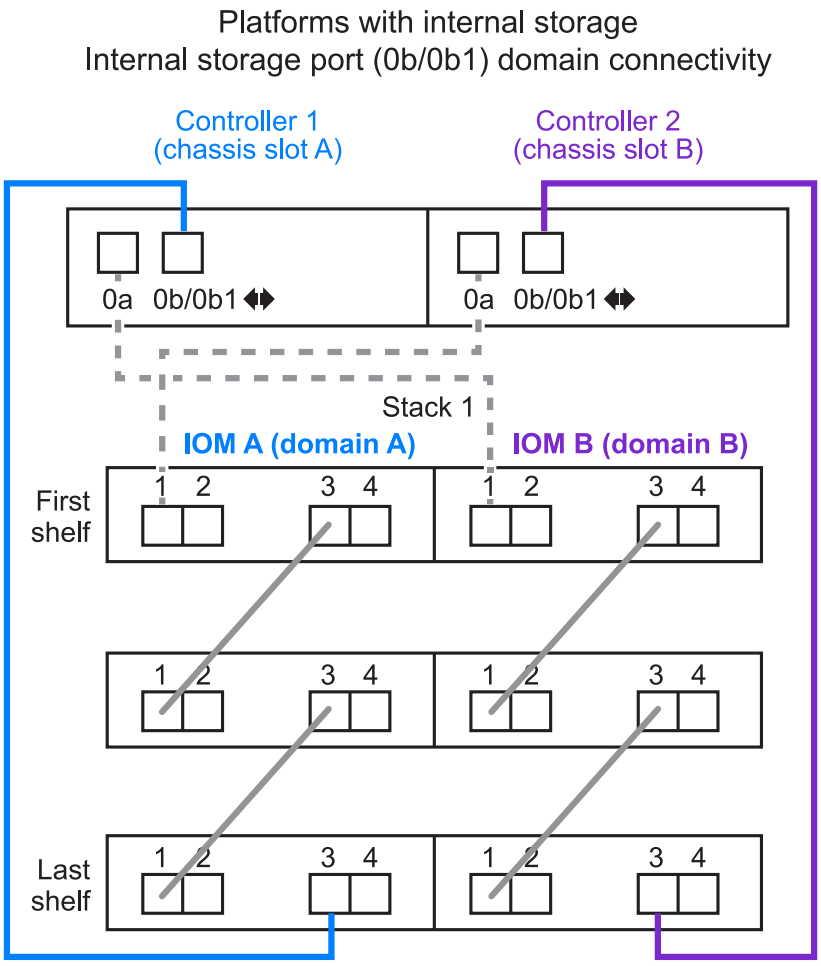

Tri-path HA connectivity

Tri-path HA connectivity is available on FAS2820 HA pairs. Tri-path HA connectivity has three paths from each controller to internal (IOM12G) and external shelves:

- Each controller's internal connection of port 0b to its local IOM12G and port 0c to its partner's IOM12G gives an HA pair multipath HA connectivity.
- The cabling of each controller's external storage ports, 0a and 0b1, gives an HA pair tri-path HA connectivity.

  Ports 0a and 0b1 are cabled across the two controllers when there are no external shelves, or they are cabled to external shelves to achieve tri-path HA connectivity.

The following shows the controller's internal connections and external cabling that achieves tri-path HA connectivity:

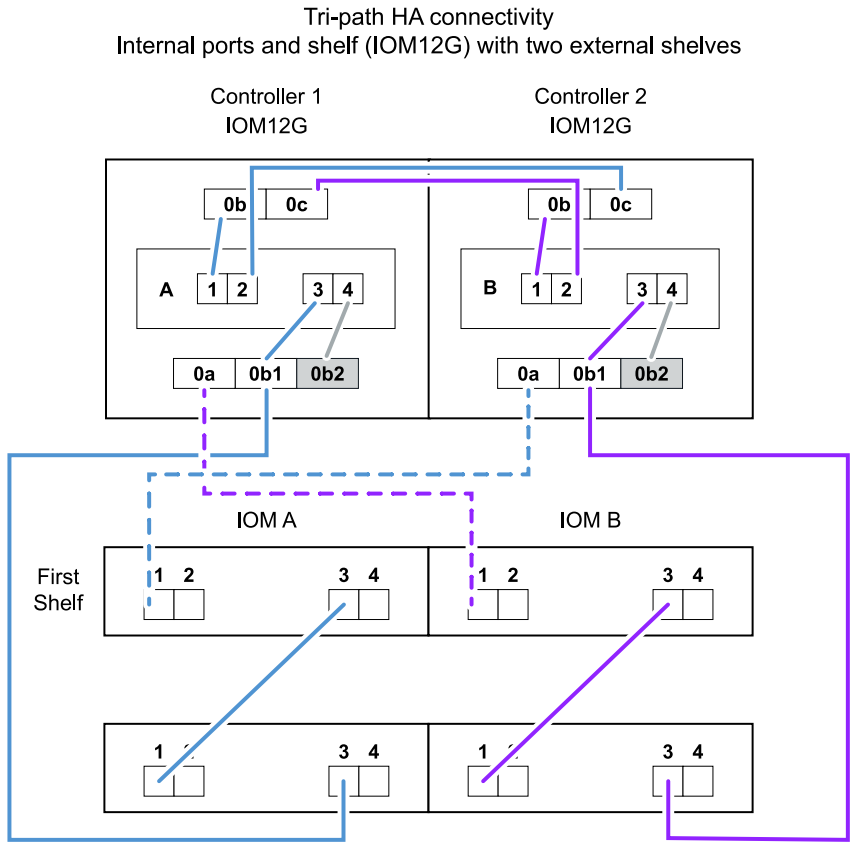

The FAS2820 external SAS ports:

- The 0a port is from the internal HBA (like other platforms with an internal shelf).
- The 0b1 port is from the internal shelf (like the 0b ports on other platforms with an internal shelf).
- The 0b2 port is not used. It is disabled. If a cable is connected to it, an error message is generated.

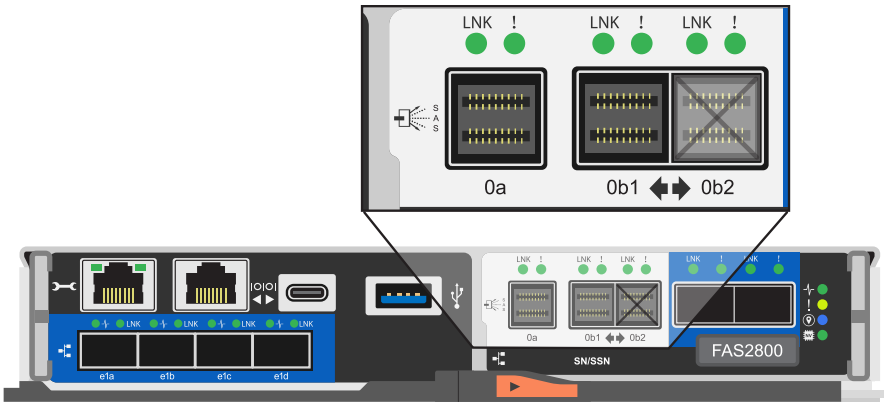

FAS2820 HA pair cabling examples can be found in the Controller-to-stack cabling worksheets and cabling examples for platforms with internal storage section.

Mini-SAS HD SAS optical cable rules

You can use mini-SAS HD SAS optical cables—multimode active optical cable (AOC) cables with mini-SAS HD-to-mini-SAS HD connectors and multimode (OM4) breakout cables with mini-SAS HD-to-LC connectors—to achieve long distance SAS connectivity for certain configurations that have disk shelves with IOM12

modules.

- Your platform and version of ONTAP must support the use of mini-SAS HD SAS optical cables: multimode active optical cable (AOC) cables with mini-SAS HD-to-mini-SAS HD connectors and multimode (OM4) breakout cables with mini-SAS HD-to-LC connectors.

  NetApp Hardware Universe

- SAS optical multimode AOC cables with mini-SAS HD-to-mini-SAS HD connectors can be used for controller-to-stack and shelf-to-shelf connections, and are available in lengths up to 50 meters.
- If you are using SAS optical multimode (OM4) breakout cables with mini-SAS HD-to-LC connectors (for patch panels), the following rules apply:
  - You can use these cables for controller-to-stack and shelf-to-shelf connections.

    If you use multimode breakout cables for shelf-to-shelf connections, you can only use them once within a stack of disk shelves. You must use multimode AOC cables to connect the remaining shelf-to-shelf connections.

    For quad-path HA and quad-path configurations, if you use multimode breakout cables for the shelf-to-shelf double-wide connections between two disk shelves, the best practice is to use identically paired breakout cables.

  - You must connect all eight (four pairs) of the LC breakout connectors to the patch panel.
  - You need to supply the patch panels and inter-panel cables.

    The inter-panel cables must be the same mode as the breakout cable: OM4 multimode.

  - Up to one pair of patch panels can be used in a path.
  - The point-to-point (mini-SAS HD-to-mini-SAS HD) path of any multimode cable cannot exceed 100 meters.

    The path includes the set of breakout cables, patch panels, and inter-panel cables.

  - The total end-to-end path (sum of point-to-point paths from the controller to the last shelf) cannot exceed 300 meters.

    The total path includes the set of breakout cables, patch panels, and inter-panel cables.

- The SAS cables can be SAS copper, SAS optical, or a mix.

  If you are using a mix of SAS copper cables and SAS optical cables, the following rules apply:

  - Shelf-to-shelf connections in a stack must be all SAS copper cables or all SAS optical cables.
  - If the shelf-to-shelf connections are SAS optical cables, the controller-to-stack connections to that stack must also be SAS optical cables.
  - If the shelf-to-shelf connections are SAS copper cables, the controller-to-stack connections to that stack can be SAS optical cables or SAS copper cables.

Cabling worksheets for multipath HA configurations - DS212C, DS224C, or DS460C

You can use the controller-to-stack cabling worksheets and cabling examples to cable your HA pair as a multipath HA configuration. This applies to shelves with IOM12/IOM12B

modules.

This information applies to platforms without internal storage.

- If needed, you can refer to SAS cabling rules and concepts for information about supported configurations, the controller slot numbering convention, shelf-to-shelf connectivity, and controller-to-shelf connectivity (including the use of port pairs).
- If needed, you can refer to How to read a worksheet to cable controller-to-stack connections for multipathed connectivity.
- Cabling examples show controller-to-stack cables as solid or dashed to distinguish controller A and C port connections from controller B and D port connections.

Controller-to-Stack Cable Type Key	
Cable Type	Description
	▪ Connects controller **A** and **C** ports to the logical **first** disk shelf in a stack ▪ The **primary** path from a controller to a stack
	▪ Connects controller **B** and **D** ports to the logical **last** disk shelf in a stack ▪ The **secondary** path from a controller to a stack

- Cables in the cabling examples and their corresponding port pairs in the worksheets are color-coded to distinguish connectivity to each stack in the HA pair.

Controller-to-Stack Cable Color Key			
Cable Color		Connects to...	From...
	Dark blue	Stack 1	Each controller by a unique port pair
	Orange	Stack 2	
	Green	Stack 3	
	Light blue	Stack 4	

- Worksheets and cabling examples show cabling port pairs in the order in which they are listed in the worksheet.

Controller-to-stack cabling worksheets and cabling examples for multipath HA configurations with quad-port SAS HBAs

You can use the completed controller-to-stack cabling worksheets and cabling examples to cable common multipath HA configurations that have quad-port SAS HBAs. These controllers do not have onboard SAS ports.

Multipath HA with one quad-port SAS HBA and one single-shelf stack

The following worksheet and cabling example uses port pair 1a/1d:

Controller-to-Stack Cabling Worksheet for Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	A	1	1a	1c				
	2	First	B	1						
B and D					~~1b~~	~~1d~~				
	1	Last	B	3	1d	1b				
	2	Last	A	3						

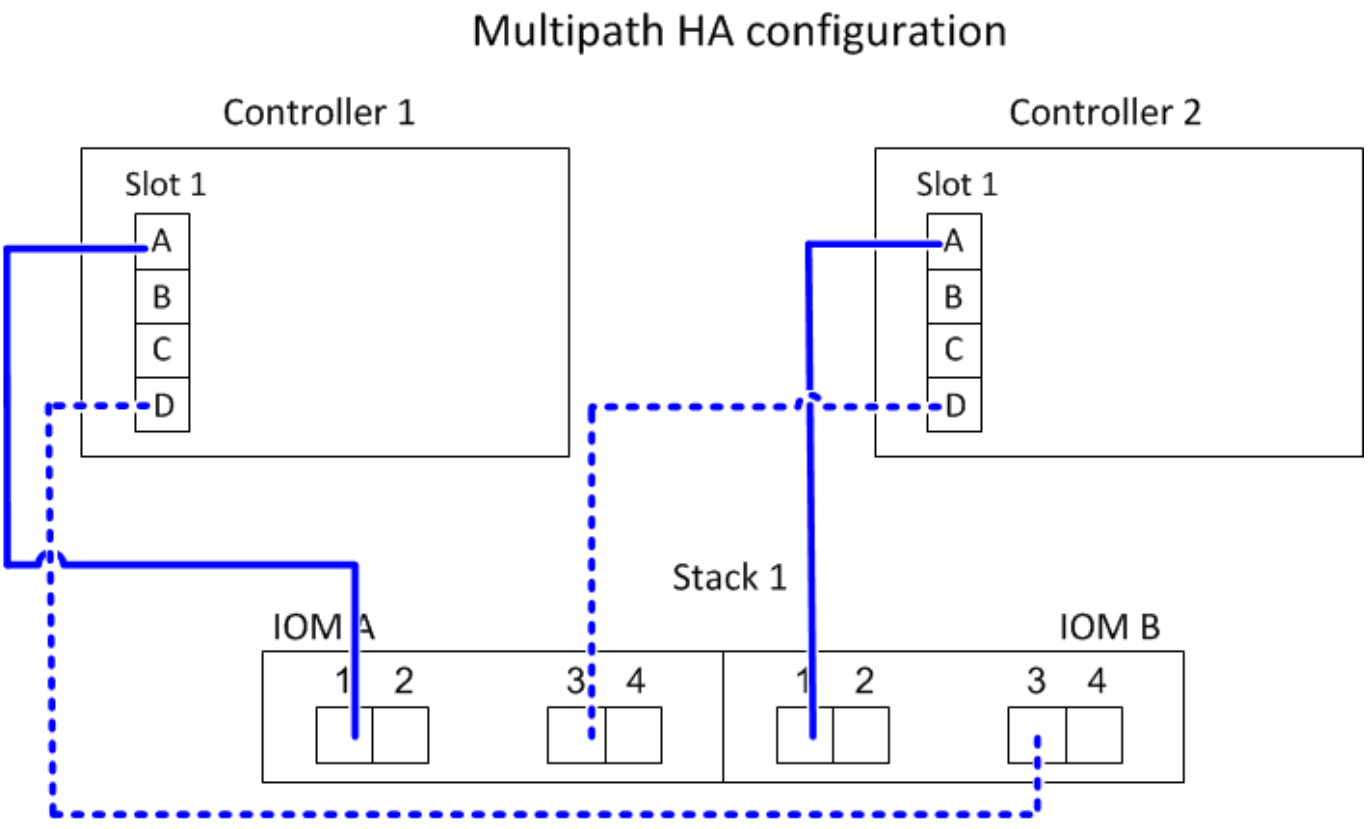

Multipath HA with one quad-port SAS HBA and two single-shelf stacks

The following worksheet and cabling example uses port pairs 1a/1d and 1c/1b:

Controller-to-Stack Cabling Worksheet for Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	A	1	1a	1c				
	2	First	B	1						
B and D					~~1b~~	~~1d~~				
	1	Last	B	3	1d	1b				
	2	Last	A	3						

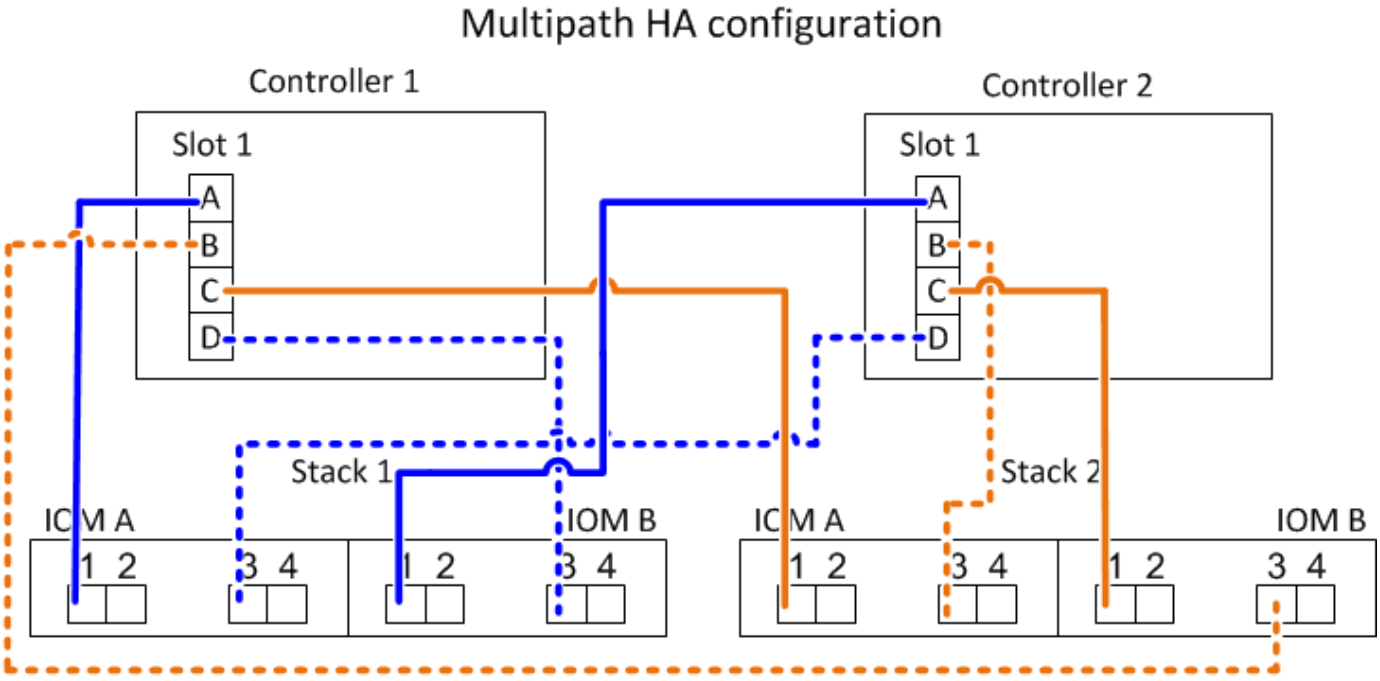

Multipath HA with two quad-port SAS HBAs and two multi-shelf stacks

Four port pairs are available for this configuration: 1a/2b, 2a/1d, 1c/2d, and 2c/1b. You can cable port pairs in the order in which they are identified (listed in the worksheet) or you can cable every other port pair (skip port pairs).

When you have more port pairs than you need to cable the stacks in your system, the best practice is to skip port pairs to optimize the SAS ports on your system. By optimizing SAS ports, you optimize your system's performance.

The following worksheet and cabling example shows port pairs being used in the order in which they are listed in the worksheet: 1a/2b, 2a/1d, 1c/2d, and 2c/1b.

Controller-to-Stack Cabling Worksheet for Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6	
		Shelf	IOM	Port	Port pairs						
A and C	1	First	A	1	1a	2a	1c	2c			
	2	First	B	1							
B and D					~~1b~~	~~2b~~	~~1d~~	~~2d~~			
	1	Last	B	3	2b	1d	2d	1b			
	2	Last	A	3							

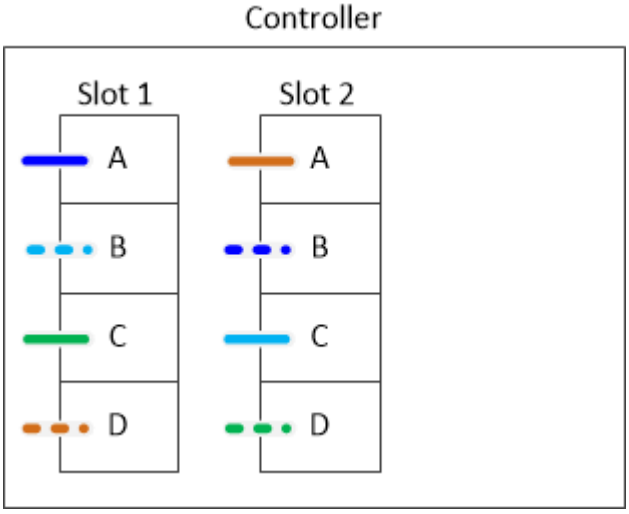

The following worksheet and cabling example shows port pairs being skipped to use every other one in the list: 1a/2b and 1c/2d.

If a third stack is added later, you use the port pair that was skipped.

Controller-to-Stack Cabling Worksheet for Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	3 ~~2~~	2 ~~3~~	4	5	6	
		Shelf	IOM	Port	Port pairs						
A and C	1	First	A	1	1a	2a	1c	2c			
	2	First	B	1							
B and D					~~1b~~	~~2b~~	~~1d~~	~~2d~~			
	1	Last	B	3	2b	1d	2d	1b			
	2	Last	A	3							

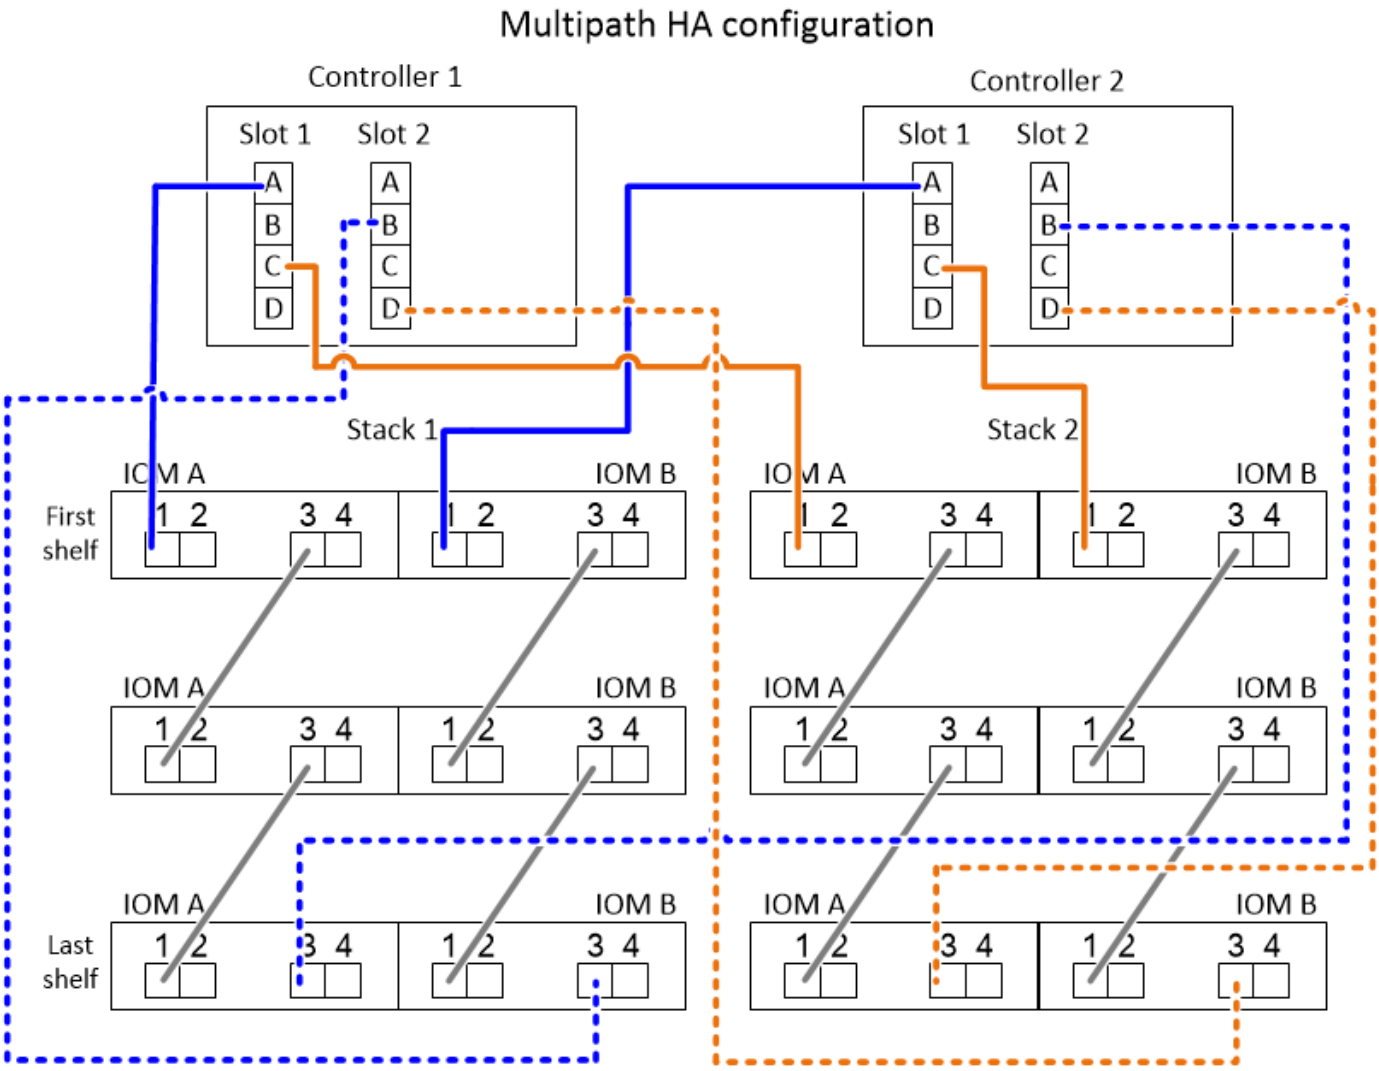

Controller-to-stack cabling worksheets and cabling examples for multipath HA configurations with four onboard SAS ports

You can use the completed controller-to-stack cabling worksheets and cabling examples to cable common multipath HA configurations that have four onboard SAS ports.

Multipath HA with four onboard SAS ports and one single-shelf stack

The following worksheet and cabling example uses port pair 0a/0d:

Controller-to-Stack Cabling Worksheet for Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	A	1	0a	0c				
	2	First	B	1						
B and D					~~0b~~	~~0d~~				
	1	Last	B	3	0d	0b				
	2	Last	A	3						

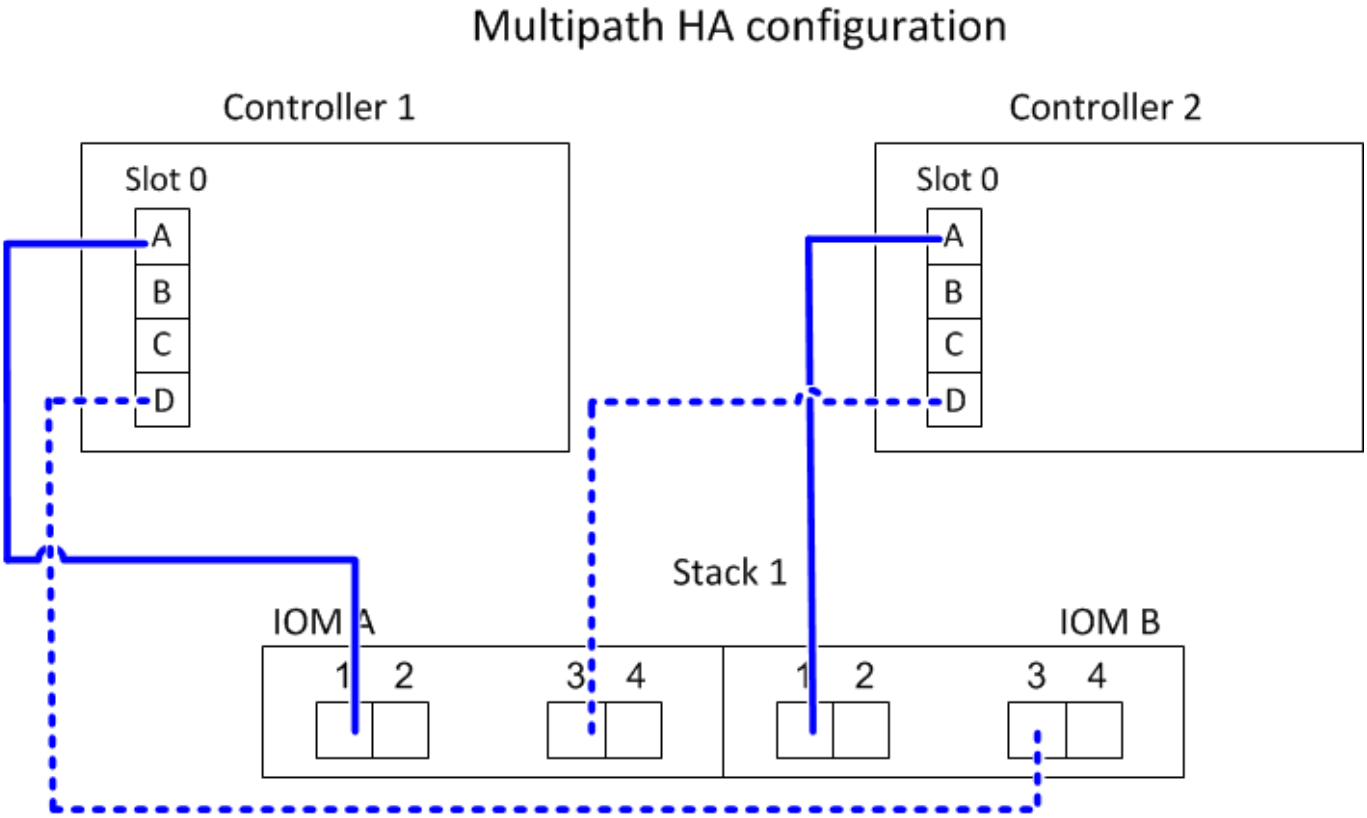

Multipath HA with four onboard SAS ports and two single-shelf stacks

The following worksheet and cabling example uses port pairs 0a/0d and 0c/0b:

Controller-to-Stack Cabling Worksheet for Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	A	1	0a	0c				
	2	First	B	1						
B and D					~~0b~~	~~0d~~				
	1	Last	B	3	0d	0b				
	2	Last	A	3						

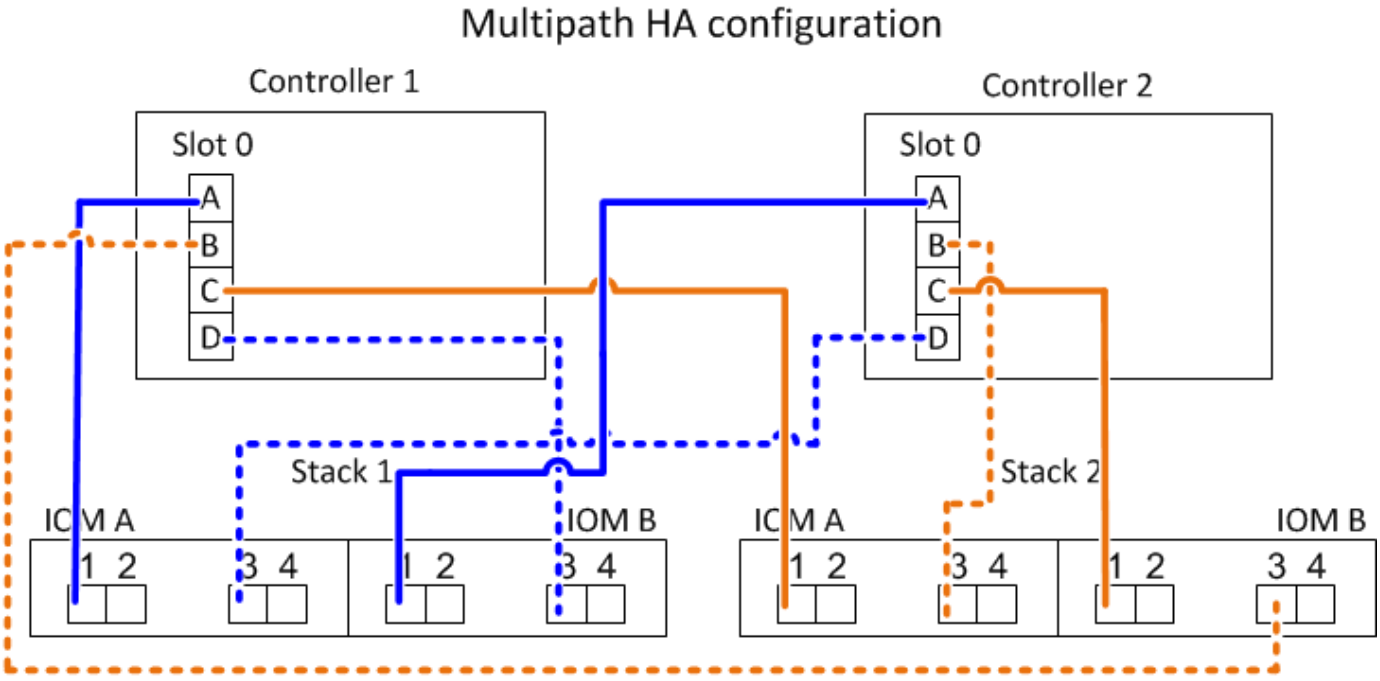

Multipath HA with four onboard SAS ports, a quad-port SAS HBA, and two multi-shelf stacks

Four port pairs are available for this configuration: 0a/1b, 1a/0d, 0c/1d, and 1c/0b. You can cable port pairs in the order in which they are identified (listed in the worksheet) or you can cable every other port pair (skip port pairs).

When you have more port pairs than you need to cable the stacks in your system, the best practice is to skip port pairs to optimize the SAS ports on your system. By optimizing SAS ports, you optimize your system's performance.

The following worksheet and cabling example shows port pairs being used in the order in which they are listed in the worksheet: 0a/1b, 1a/0d, 0c/1d, and 1c/0b.

Controller-to-Stack Cabling Worksheet for Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	A	1	0a	1a	0c	1c		
	2	First	B	1						
B and D					~~0b~~	~~1b~~	~~0d~~	~~1d~~		
	1	Last	B	3	1b	0d	1d	0b		
	2	Last	A	3						

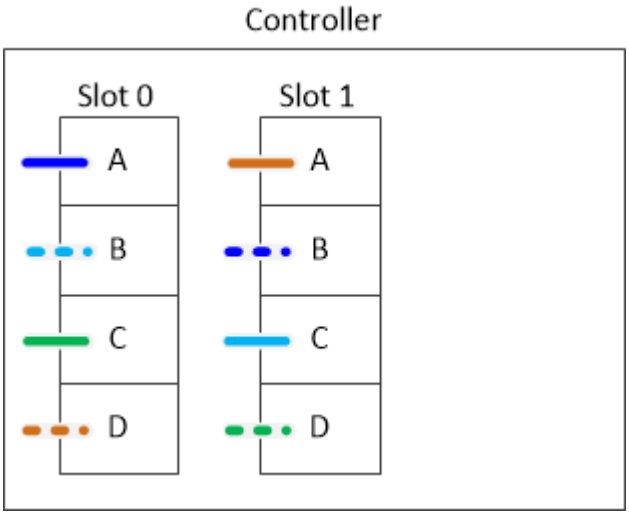

The following worksheet and cabling example shows port pairs being skipped to use every other one in the list: 0a/1b and 0c/1d.

If a third stack is added later, you use the port pair that was skipped.

Controller-to-Stack Cabling Worksheet for Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	3 ~~2~~	2 ~~3~~	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	A	1	0a	1a	0c	1c		
	2	First	B	1						
B and D					~~0b~~	~~1b~~	~~0d~~	~~1d~~		
	1	Last	B	3	1b	0d	1d	0b		
	2	Last	A	3						

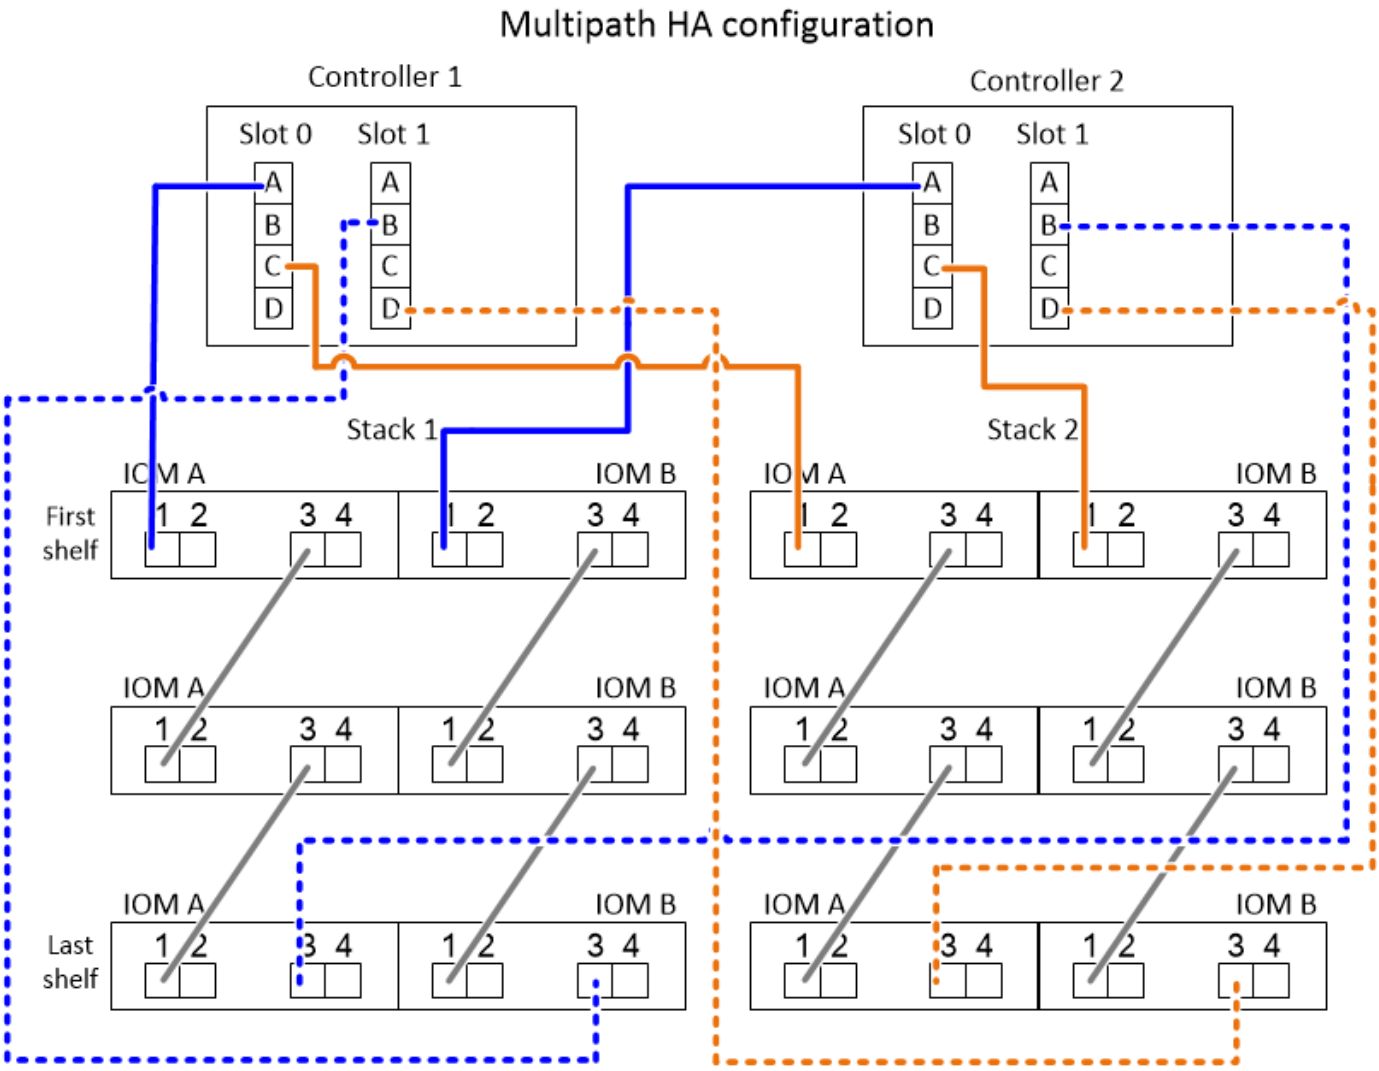

Cabling worksheets for internal storage - DS212C, DS224C, or DS460C

You can use the completed controller-to-stack cabling worksheets and cabling examples to cable platforms with internal storage. This applies to shelves with IOM12/IOM12B modules.

This information does not apply to FAS25XX platforms.

- If needed, you can refer to SAS cabling rules and concepts for information about supported configurations, shelf-to-shelf connectivity, and controller-to-shelf connectivity.
- Cabling examples show controller-to-stack cables as solid or dashed to distinguish controller 0b/0b1 port connections from controller 0a port connections.

Controller-to-stack cable type key: AFF and FAS platforms with onboard storage (except FAS25XX)	
Cable Type	**Description**
———	• Connects controller **0b** or **0b1** port to the logical **last** disk shelf in the stack • The **primary** path from a controller to the stack The internal storage connection
— — — —	• Connects controller **0a** port to the logical **first** disk shelf in the stack • The **secondary** path from a controller to the stack The internal HBA connection

- Cabling examples show controller-to-stack connections and shelf-to-shelf connections in two different colors to distinguish connectivity through IOM A (domain A) and IOM B (domain B).

Cable color key: AFF and FAS platforms with onboard storage (except FAS25XX)		
Cable Color		**Connects...**
	Light blue	IOM A (domain A)
	Purple	IOM B (domain B)

FAS2820 platform in a multipath HA configuration with no external shelves

The following example shows that no cabling is needed to acheive multipath HA connectivity:

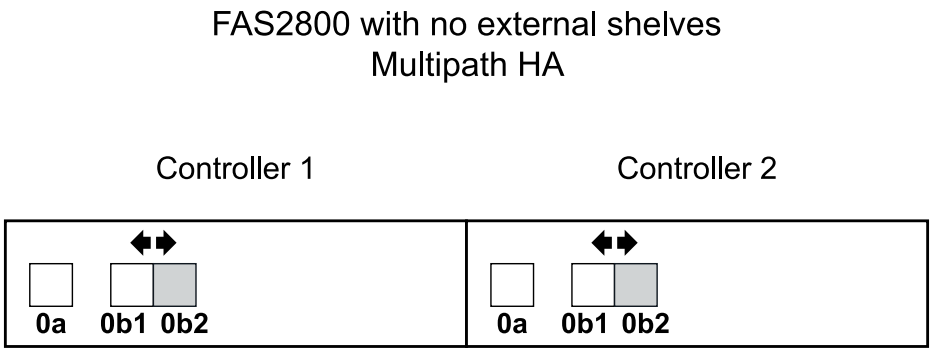

FAS2820 platform in a tri-path HA configuration with no external shelves

The following cabling example shows required cabling between the two controllers to achieve tri-path connectivity:

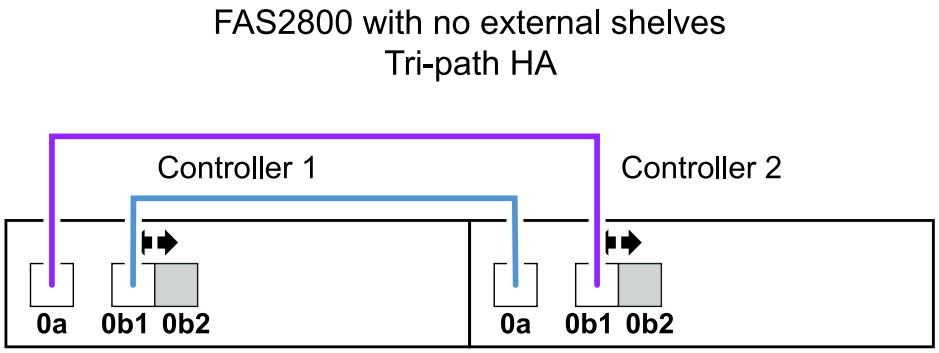

FAS2820 platform in a tri-path HA configuration with one multi-shelf stack

The following worksheet and cabling example uses port pair 0a/0b1:

Controller-to-stack cabling worksheet: FAS2800 platform											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	B	1	0a					
	2	First	A	1						
B and D	1	Last	A	3	0b1					
	2	Last	B	3						

FAS2800 platform
Tri-path HA configuration

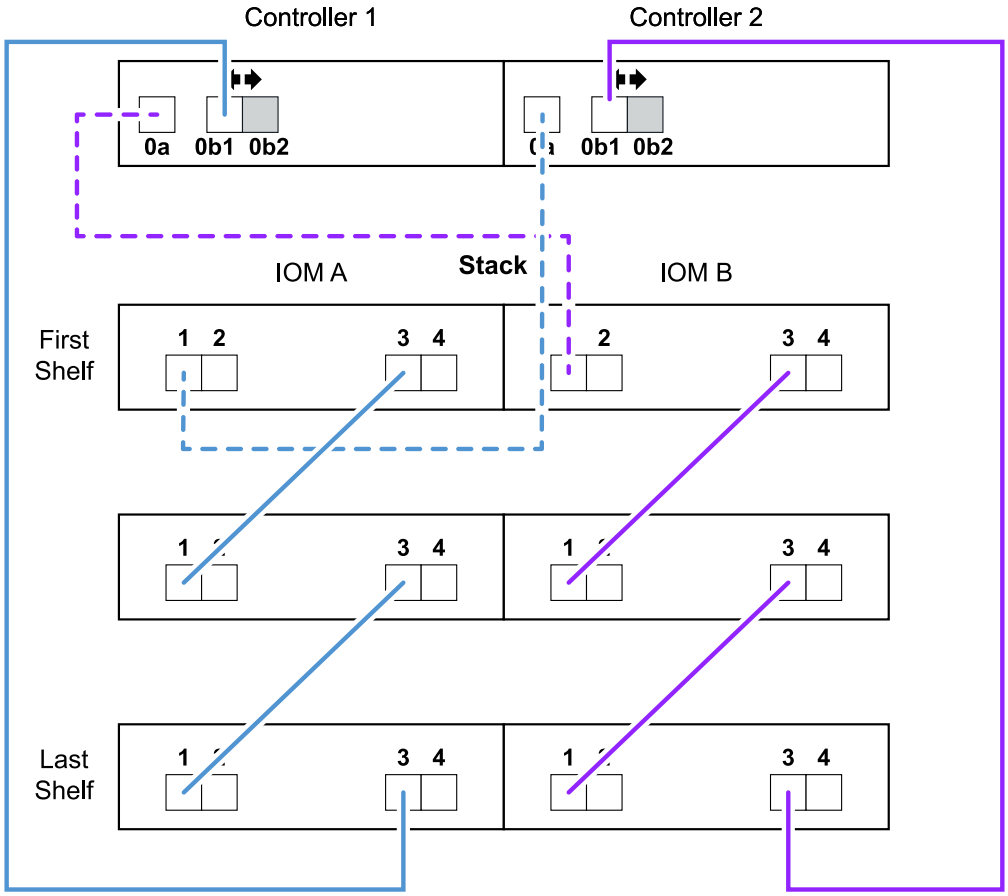

Platforms with internal storage in a multipath HA configuration with one multi-shelf stack

The following worksheet and cabling example uses port pair 0a/0b:

This section does not apply to FAS2820 or FAS25XX systems.

Controller-to-stack cabling worksheet: AFF and FAS platforms with onboard storage											
Controller SAS ports	**Controllers**	**Cable to disk shelf IOMs**			**Stacks**						
					1	2	3	4	5	6
		Shelf	**IOM**	**Port**		**Port pairs**				
A and C	1	First	B	1	0a					
	2	First	A	1						
B and D	1	Last	A	3	0b					
	2	Last	B	3						

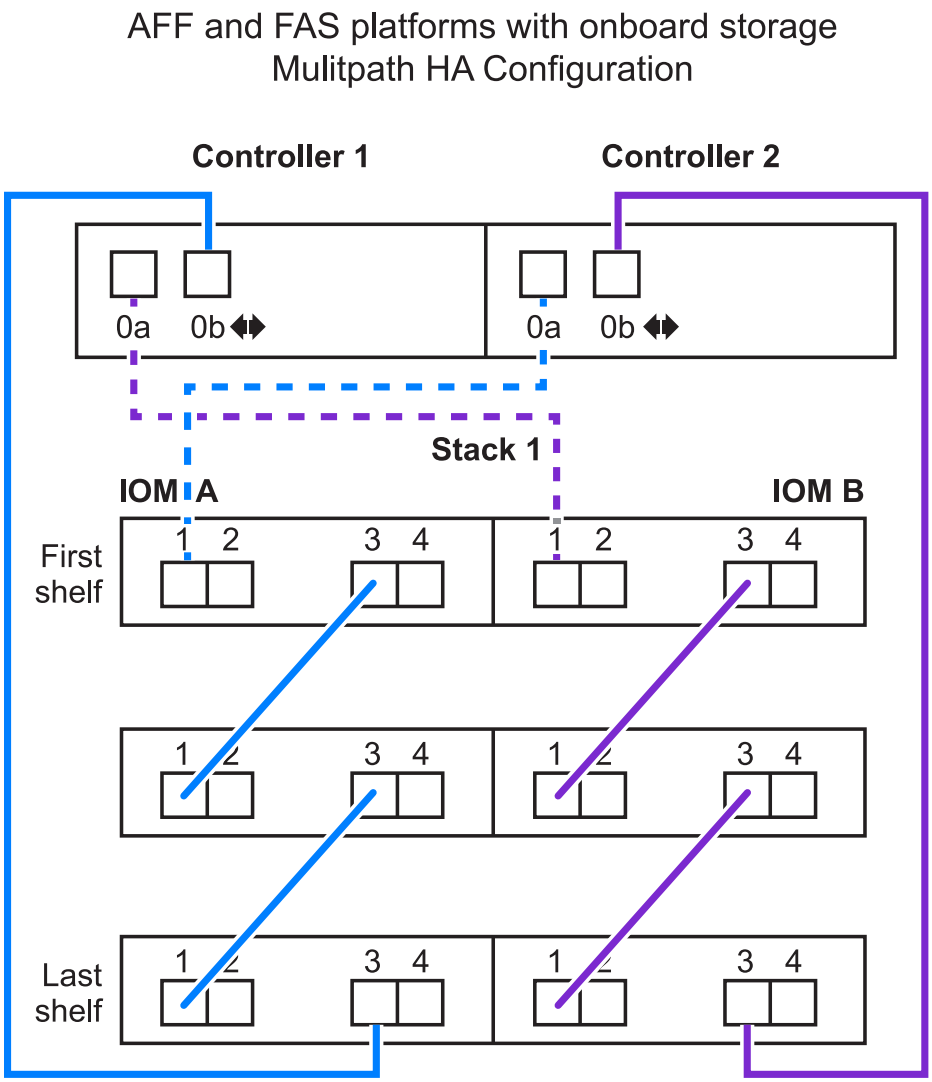

FAS2600 series multipath configuration with one multi-shelf stack

The following worksheets and cabling examples use port pair 0a/0b.

In this example, the controller is installed in slot A of the chassis. When a controller is located in slot A of the chassis, its internal storage port (0b) is in domain A (IOM A); therefore, port 0b must connect to domain A (IOM A) in the stack.

Controller-to-Stack Cabling Worksheet (FAS2600 series)											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6	
		Shelf	IOM	Port		Port pairs					
A and C	1	First	B	1	0a						
	~~2~~	~~First~~	~~A~~	~~1~~							
B and D	1	Last	A	3	0b						
	~~2~~	~~Last~~	~~B~~	~~3~~							

FAS2600 series multipath configuration

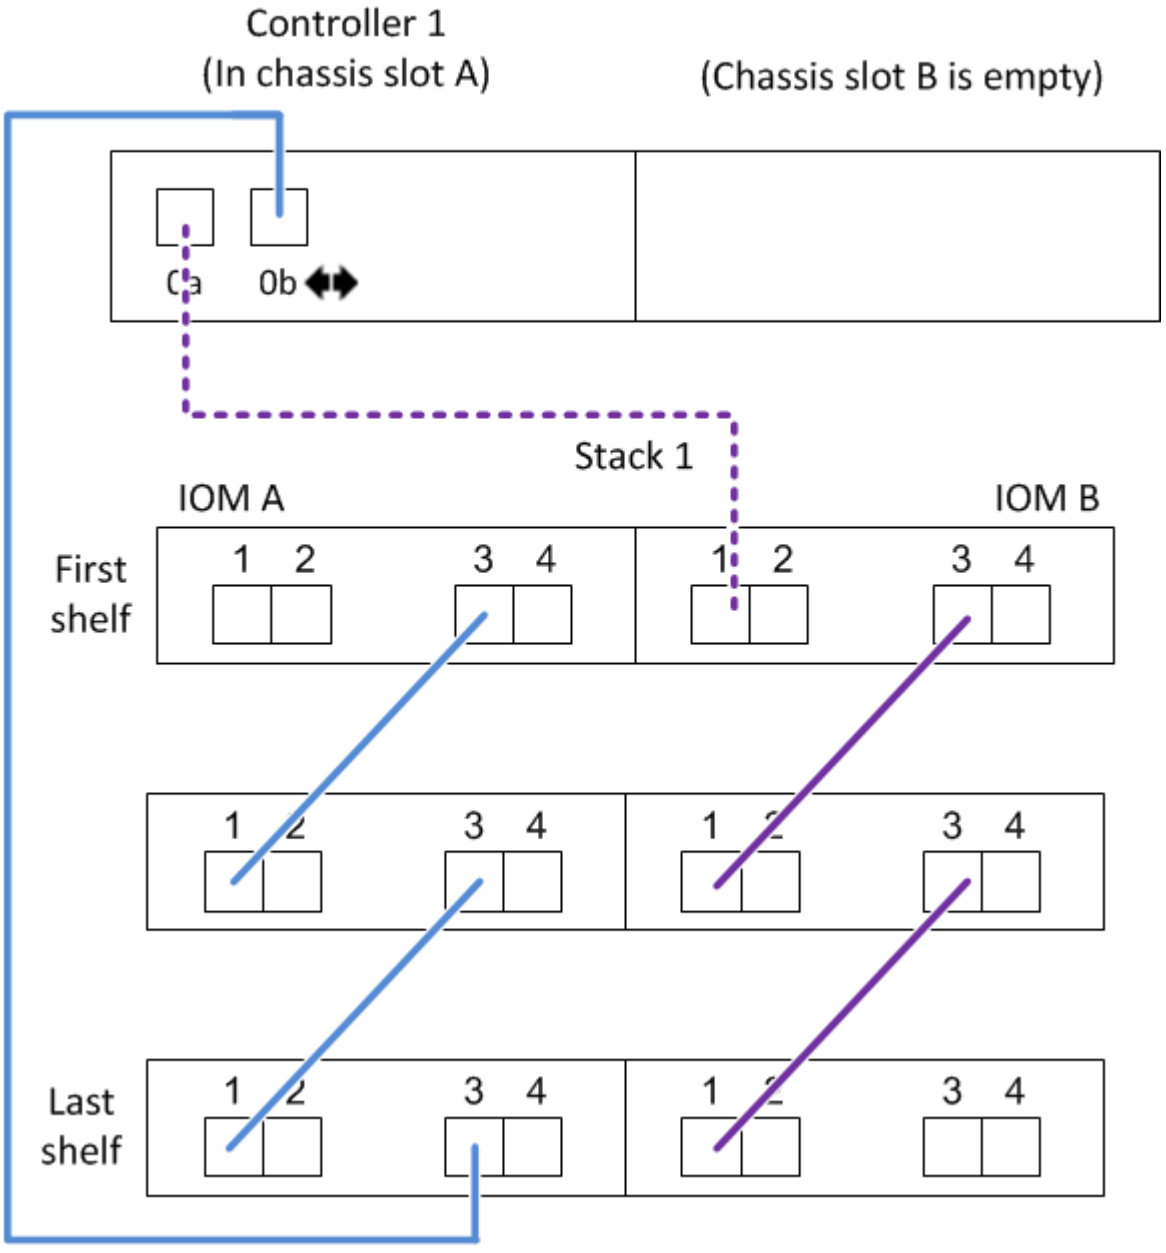

In this example, the controller is installed in slot B of the chassis. When a controller is located in slot B of the chassis, its internal storage port (0b) is in domain B (IOM B); therefore, port 0b must connect to domain B (IOM B) in the stack.

Controller-to-Stack Cabling Worksheet (FAS2600 series)											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6
		Shelf	IOM	Port		Port pairs				
A and C	1	First	B	1	0a					
	~~2~~	~~First~~	~~A~~	~~1~~						
B and D	1	Last	A	3	0b					
	~~2~~	~~Last~~	~~B~~	~~3~~						

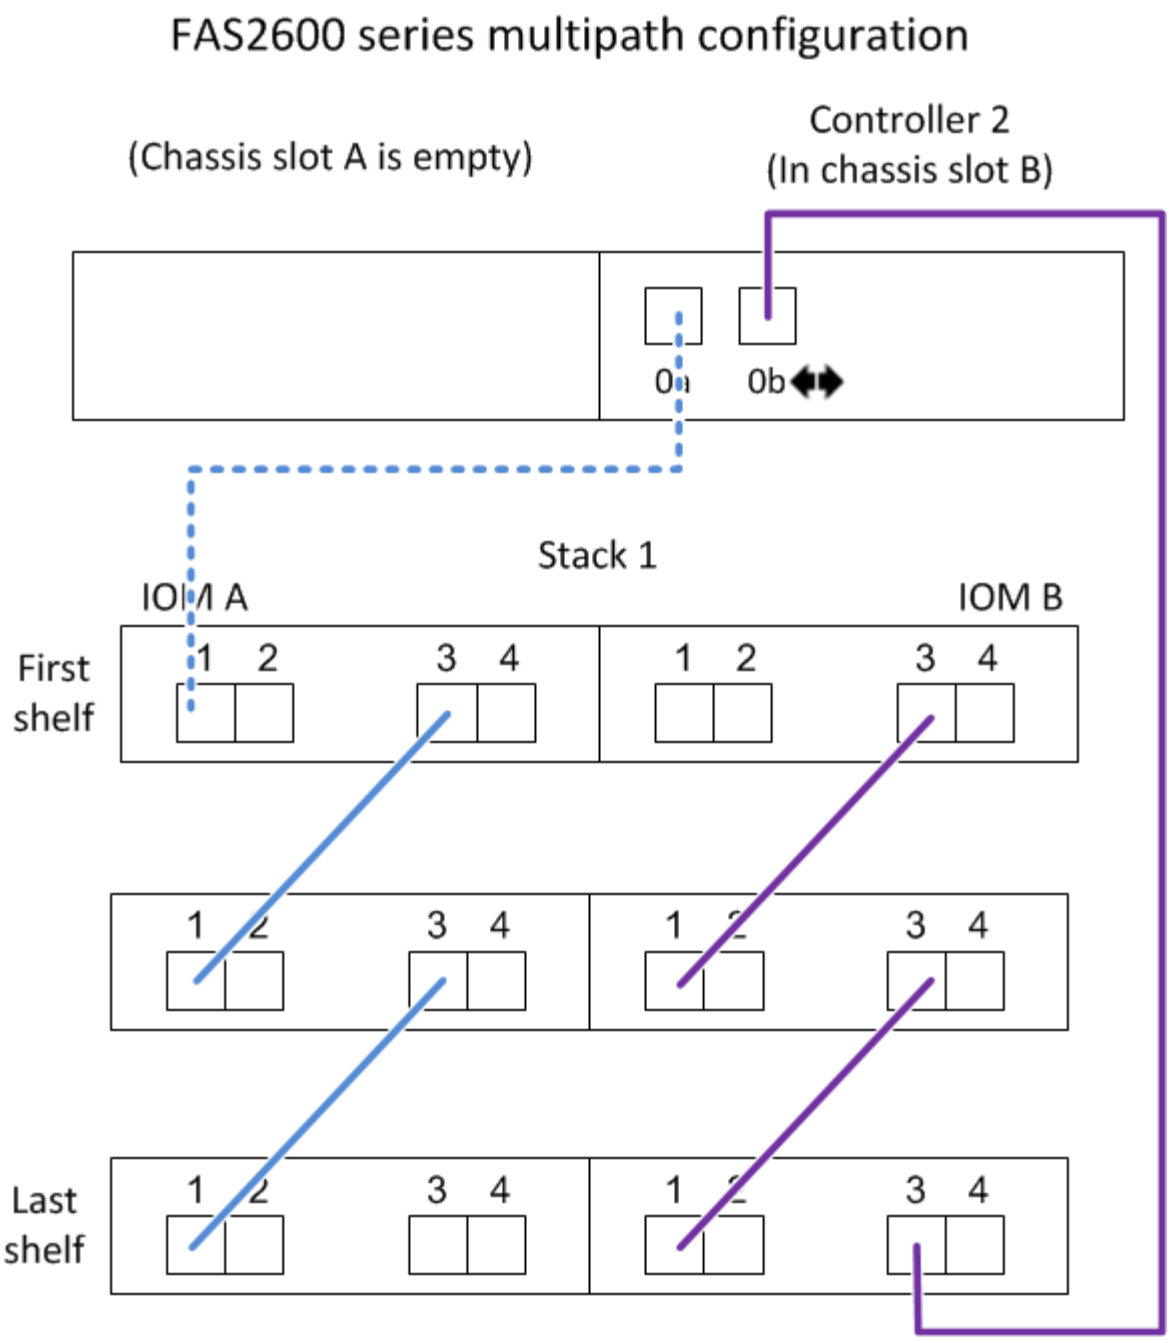

Cabling worksheet for a quad-path HA configuration with two quad-port SAS HBAs - DS212C, DS224C, or DS460C

You can use the completed controller-to-stack cabling worksheet and cabling example to cable a quad-path HA configuration that has two quad-port SAS HBAs. This applies to shelves with IOM12/IOM12B modules.

- If needed, you can refer to SAS cabling rules for information about supported configurations, the controller slot numbering convention, shelf-to-shelf connectivity, and controller-to-shelf connectivity (including the use of port pairs).
- If needed, you can refer to How to read a worksheet to cable controller-to-stack connections for quad-pathed connectivity.
- The cabling example shows controller-to-stack cables as solid or dashed to distinguish controller A and C port connections from controller B and D port connections.

Controller-to-Stack Cable Type Key	
Cable Type	Description
	▪ Connects controller **A** and **C** ports to the logical **first** disk shelf in a stack ▪ The **primary** path from a controller to a stack
	▪ Connects controller **B** and **D** ports to the logical **last** disk shelf in a stack ▪ The **secondary** path from a controller to a stack

- Cables in the cabling examples and their corresponding port pairs in the worksheets are color-coded to distinguish connectivity to each stack in the HA pair.

Controller-to-Stack Cable Color Key			
Cable Color		Connects to...	From...
	Dark blue	Stack 1	Each controller by a unique port pair
	Orange	Stack 2	

- The cabling example visually distinguishes the two sets of multipathed cabling needed to achieve quad-pathed connectivity for each controller to each stack in an HA pair or single-controller configuration.

  The first set of multipathed cabling is referred to as “multipathed”. The second set of multipathed cabling is referred to as “quad-pathed”. The second set of cabling is referred to as “quad-pathed” because completing this set of cabling gives you the quad-pathed connectivity.

Controller-to-Stack Quad-Pathed Connectivity Key			
Quad-pathed connectivity consists of two sets of cabling		Shown by color-coded ports on controllers and IOMs	Description
Set 1	Multipathed	No color	Ports (on controllers and IOMs) cabled with multipathed connectivity are shown without a color.
Set 2	Quad-pathed	The cable color associated with the applicable stack	Ports (on controllers and IOMs) cabled with quad-pathed connectivity are the same color as the cables connecting the stack, as shown in the “Controller-to-Stack Cable Color Key”.

- The worksheet example shows port pairs designated for multipathed cabling or quad-pathed cabling to the applicable stack.

  Each port pair designated for multipathed cabling is encircled by an oval that is the color associated with the stack it is cabled to. Each port pair designated for quad-pathed cabling is encircled by a rectangle that is the color associated with the stack it is cabled to.

Quad-path HA with two quad-port SAS HBAs and two multi-shelf stacks

The following worksheet and cabling example uses port pairs 1a/2b (multipathed) and 2a/1d (quad-pathed) for stack 1, and port pairs 1c/2d (multipathed) and 2c/1b (quad-pathed) for stack2.

Controller-to-Stack Cabling Worksheet for Quad-Pathed Connectivity									
Controller SAS ports	Controllers	Cable to disk shelf IOMs				Stacks			
		Shelf	IOM	Port		1		2	
				Multipathed	Quad-pathed	Port pairs			
A and C	1	First	A	1	2	1a	2a	1c	2c
	2	First	B	1	2				
B and D						~~1b~~	~~2b~~	~~1d~~	~~2d~~
	1	Last	B	3	4	2b	1d	2d	1b
	2	Last	A	3	4				

Quad-path HA configuration

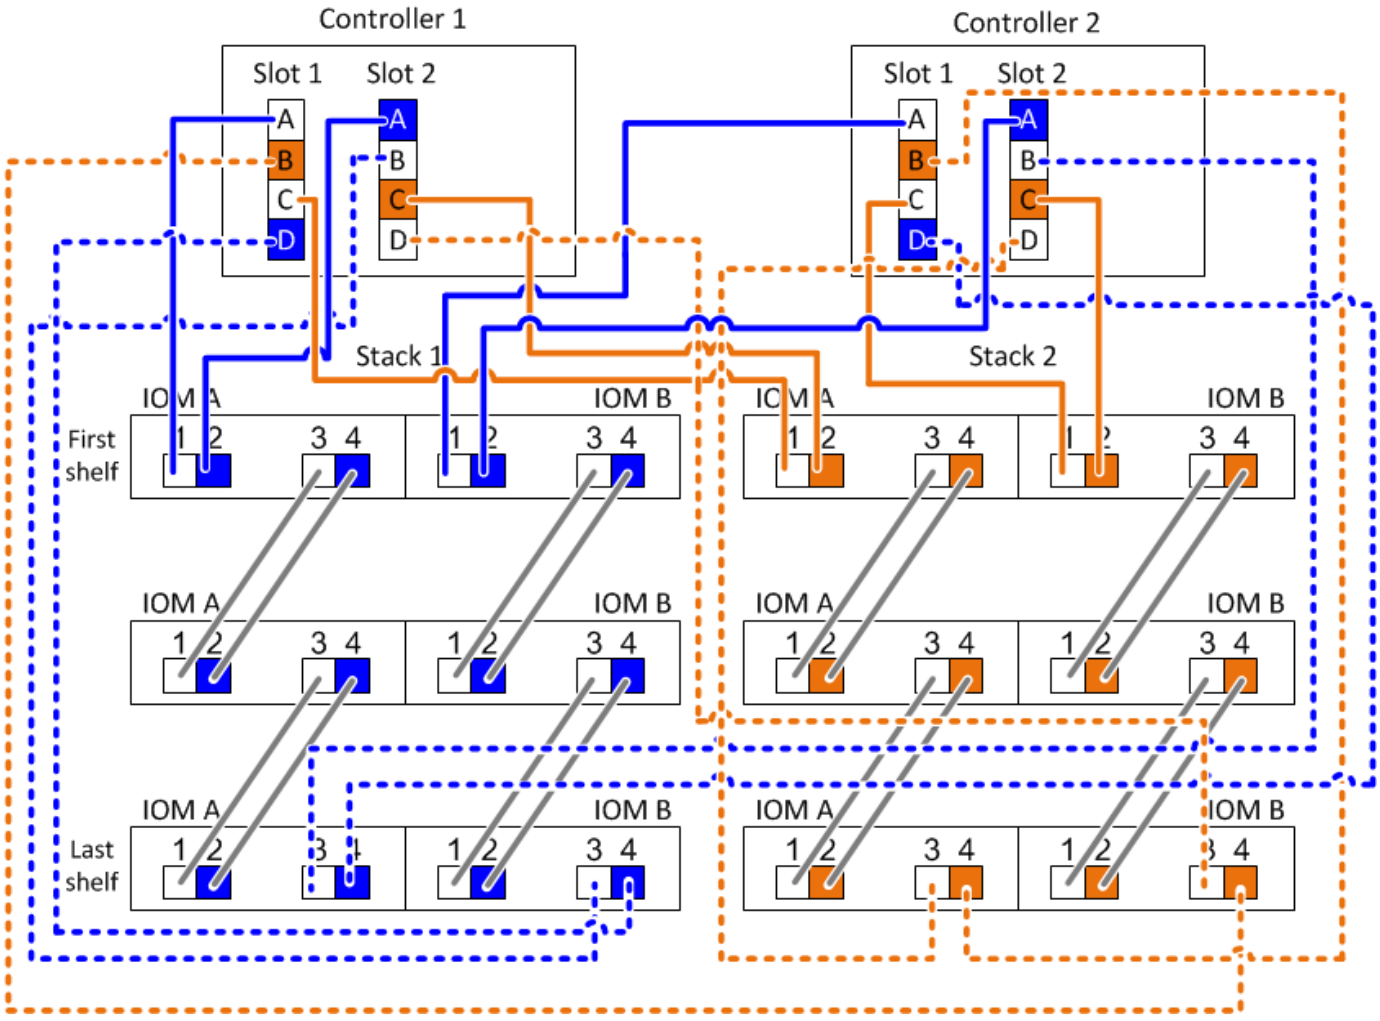

Cabling worksheet for multipathed connectivity - DS212C, DS224C, or DS460C

By completing the worksheet template, you can define the controller SAS port pairs you can use to cable controllers to stacks of disk shelves with IOM12/IOM12B modules to achieve multipathed connectivity in an HA pair or single-controller configuration. You can also use the completed worksheet to walk yourself through cabling the multipathed connections for your configuration.

Before you begin

If you have a platform with internal storage, use the following worksheet:

Controller-to-stack cabling worksheets and cabling examples for platforms with internal storage

About this task

- This procedure and worksheet template is applicable to cabling multipathed connectivity for a multipath HA or multipath configuration with one or more stacks.

  Examples of completed worksheets are provided for multipath HA and multipath configurations.

  A configuration with two quad-port SAS HBAs and two stacks of disk shelves with IOM12/IOM12B modules is used for the worksheet examples.

- The worksheet template allows for up to six stacks; you need to add more columns if needed.
- If needed, you can refer to the SAS cabling rules and concepts for information about supported configurations, the controller slot numbering convention, shelf-to-shelf connectivity, and controller-to-shelf connectivity (including use of port pairs).
- If needed, after you complete the worksheet, you can refer to How to read a worksheet to cable controller-to-stack connections for multipathed connectivity

Controller-to-Stack Cabling Worksheet Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	A	1						
	2	First	B	1						
B and D										
	1	Last	B	3						
	2	Last	A	3						

Steps

1. In the boxes above the gray boxes, list all SAS A ports on your system, and then all SAS C ports on your system in sequence of slots (0, 1, 2, 3, and so on).

   For example: 1a, 2a, 1c, 2c

2. In the gray boxes, list all SAS B ports on your system, and then all SAS D ports on your system in sequence of slots (0, 1, 2, 3 and so on).

   For example: 1b, 2b, 1d, 2d

3. In the boxes below the gray boxes, rewrite the D and B port list so that the first port in the list is moved to the end of the list.

   For example: 2b, 1d, 2d, 1b

4. Circle (designate) a port pair for each stack.

   When all port pairs are being used to cable the stacks in your system, circle port pairs in the order in which they are defined (listed) in the worksheet.

   For example, in a multipath HA configuration with eight SAS ports and four stacks, port pair 1a/2b is cabled to stack 1, port pair 2a/1d is cabled to stack 2, port pair 1c/2d is cabled to stack3, and port pair 2c/1b is cabled to stack 4.

Controller-to-Stack Cabling Worksheet for Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	2	3	4	5	6	
		Shelf	IOM	Port	Port pairs						
A and C	1	First	A	1	1a	2a	1c	2c			
	2	First	B	1							
B and D					~~1b~~	~~2b~~	~~1d~~	~~2d~~			
	1	Last	B	3	2b	1d	2d	1b			
	2	Last	A	3							

   When not all port pairs are needed to cable the stacks in your system, skip port pairs (use every other port pair).

   For example, in a multipath HA configuration with eight SAS ports and two stacks, port pair 1a/2b is cabled to stack 1 and port pair 1c/2d is cabled to stack 2. If two additional stacks are hot-added later, port pair 2a/1d is cabled to stack 3 and port pair 2c/1b is cabled to stack 4.

   When you have more port pairs than you need to cable the stacks in your system, the best practice is to skip port pairs to optimize the SAS ports on your system. By optimizing SAS ports, you optimize your system's performance.

Controller-to-Stack Cabling Worksheet Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	3 ~~2~~	2 ~~3~~	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	A	1	1a	2a	1c	2c		
	2	First	B	1						
B and D					~~1b~~	~~2b~~	~~1d~~	~~2d~~		
	1	Last	B	3	2b	1d	2d	1b		
	2	Last	A	3						

You can use your completed worksheet to cable your system.

5. If you have a single-controller (multipath) configuration, cross out the information for controller 2.

Controller-to-Stack Cabling Worksheet Multipathed Connectivity										
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks					
					1	2	3	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	A	1	1a	2a	1c	2c		
	~~2~~	~~First~~	~~B~~	~~1~~						
B and D					~~1b~~	~~2b~~	~~1d~~	~~2d~~		
	1	Last	B	3	2b	1d	2d	1b		
	~~2~~	~~Last~~	~~A~~	~~3~~						

You can use your completed worksheet to cable your system.

Cabling worksheet for quad-pathed connectivity - DS212C, DS224C, or DS460C

By completing the worksheet template, you can define the controller SAS port pairs you can use to cable controllers to stacks of disk shelves with IOM12/IOM12B modules to achieve quad-pathed connectivity in an HA pair or single-controller configuration. You can also use the completed worksheet to walk yourself through cabling the quad-pathed connections for your configuration.

About this task

- This procedure and worksheet template is applicable to cabling quad-pathed connectivity for a quad-path HA or quad-path configuration with one or more stacks.

  Examples of completed worksheets are provided for quad-path HA and quad-path configurations.

A configuration with two quad-port SAS HBAs and two stacks of disk shelves with IOM12/IOM12B modules is used for the worksheet examples.

- The worksheet template allows for up to two stacks; you need to add more columns if needed.
- Quad-pathed connectivity for controller-to-stack connections consists of two sets of multipathed cabling: the first set of cabling is referred to as “multipathed”; the second set of cabling is referred to as “quad-pathed”.

  The second set of cabling is referred to as “quad-pathed” because completing this set of cabling gives you the quad-pathed connectivity from a controller to a stack in an HA pair or single-controller configuration.

- Disk shelf IOM ports 1 and 3 are always used for multipathed cabling and IOM ports 2 and 4 are always used for quad-pathed cabling, as designated by the worksheet column headings.
- In the worksheet examples, port pairs are designated for multipathed cabling or quad-pathed cabling to the applicable stack.

  Each port pair designated for multipathed cabling is encircled by an oval that is the color associated with the stack it is cabled to. Each port pair designated for quad-pathed cabling is encircled by a rectangle that is the color associated with the stack it is cabled to. Stack 1 is associated with the color blue; stack 2 is associated with the color orange.

- If needed, you can refer to SAS cabling rules and concepts for information about the controller slot numbering convention, shelf-to-shelf connectivity, and controller-to-shelf connectivity (including the use of port pairs).
- If needed, after you complete the worksheet, you can refer to How to read a worksheet to cable controller-to-stack connections for quad-pathed connectivity.

Controller-to-Stack Cabling Worksheet for Quad-Pathed Connectivity									
Controller SAS ports	Controllers	Cable to disk shelf IOMs				Stacks			
		Shelf	IOM	Port		1		2	
				Multipathed	Quad-pathed	Port pairs			
A and C	1	First	A	1	2				
	2	First	B	1	2				
B and D									
	1	Last	B	3	4				
	2	Last	A	3	4				

Steps

1. In the boxes above the gray boxes, list all SAS A ports on your system, and then all SAS C ports on your system in sequence of slots (0, 1, 2, 3, and so on).

   For example: 1a, 2a, 1c, 2c

2. In the gray boxes, list all SAS B ports on your system, and then all SAS D ports on your system in sequence of slots (0, 1, 2, 3 and so on).

For example: 1b, 2b, 1d, 2d

3. In the boxes below the gray boxes, rewrite the D and B port list so that the first port in the list is moved to the end of the list.

   For example: 2b, 1d, 2d, 1b

4. Identify the two sets of port pairs to connect to stack 1 by drawing an oval around the first set of port pairs and a rectangle around the second set of port pairs.

   Both sets of cabling are needed to achieve quad-pathed connectivity from each controller to stack 1 in your HA pair or single-controller configuration.

   The following example uses port pair 1a/2b for the multipathed cabling and port pair 2a/1d for the quad-pathed cabling to stack 1.

Controller-to-Stack Cabling Worksheet for Quad-Pathed Connectivity									
Controller SAS ports	Controllers	Cable to disk shelf IOMs				Stacks			
		Shelf	IOM	Port		1		2	
				Multipathed	Quad-pathed	Port pairs			
A and C	1	First	A	1	2	1a	2a	1c	2c
	2	First	B	1	2				
B and D						~~1b~~	~~2b~~	~~1d~~	~~2d~~
	1	Last	B	3	4	2b	1d	2d	1b
	2	Last	A	3	4				

5. Identify the two sets of port pairs to connect to stack 2 by drawing an oval around the first set of port pairs and a rectangle around the second set of port pairs.

   Both sets of cabling are needed to achieve quad-pathed connectivity from each controller to stack 1 in your HA pair or single-controller configuration.

   The following example uses port pair 1c/2d for the multipathed cabling and port pair 2c/1b for the quad-pathed cabling to stack 2.

Controller-to-Stack Cabling Worksheet for Quad-Pathed Connectivity									
Controller SAS ports	Controllers	Cable to disk shelf IOMs				Stacks			
		Shelf	IOM	Port		1		2	
				Multipathed	Quad-pathed	Port pairs			
A and C	1	First	A	1	2	1a	2a	1c	2c
	2	First	B	1	2				
B and D						~~1b~~	~~2b~~	~~1d~~	~~2d~~
	1	Last	B	3	4	2b	1d	2d	1b
	2	Last	A	3	4				

6. If you have a quad-path (single-controller) configuration, cross out the information for controller 2; you only need controller 1 information to cable the controller-to-stack connections.

   The following example shows that the information for controller 2 is crossed out.

Controller-to-Stack Cabling Worksheet for Quad-Pathed Connectivity									
Controller SAS ports	Controllers	Cable to disk shelf IOMs				Stacks			
		Shelf	IOM	Port		1		2	
				Multipathed	Quad-pathed	Port pairs			
A and C	1	First	A	1	2	1a	2a	1c	2c
	~~2~~	~~First~~	~~B~~	~~1~~	~~2~~				
B and D						~~1b~~	~~2b~~	~~1d~~	~~2d~~
	1	Last	B	3	4	2b	1d	2d	1b
	~~2~~	~~Last~~	~~A~~	~~3~~	~~4~~				

How to read a worksheet to cable controller-to-stack connections for multipathed connectivity - DS212C, DS224C, or DS460C

You can use this example to guide you through how to read and apply a completed worksheet to cable controller-to-stack connections for disk shelves with IOM12/IOM12B modules for multipathed connectivity.

Before you begin

If you have a platform with internal storage, use the following worksheet:

Controller-to-stack cabling worksheets and cabling examples for platforms with internal storage

About this task

- This procedure references the following worksheet and cabling example to demonstrate how to read a worksheet to cable controller-to-stack connections.

The configuration used in this example is a multipath HA configuration with two quad-port SAS HBAs (eight SAS ports) on each controller and two stacks of disk shelves with IOM12/IOM12B modules. Port pairs are cabled by skipping every other port pair in the worksheet.

> When you have more port pairs than you need to cable the stacks in your system, the best practice is to skip port pairs to optimize the SAS ports on your system. By optimizing SAS ports, you optimize your system's performance.

- If you have a single-controller configuration, skip substeps b and d for cabling to a second controller.
- If needed, you can refer to SAS cabling rules and concepts for information about the controller slot numbering convention, shelf-to-shelf connectivity, and controller-to-shelf connectivity (including the use of port pairs).

The port pairs are cabled using every other port pair in the worksheet: 1a/2b and 1c/2d.

Controller-to-Stack Cabling Worksheet Multipathed Connectivity											
Controller SAS ports	Controllers	Cable to disk shelf IOMs			Stacks						
					1	3 ~~2~~	2 ~~3~~	4	5	6
		Shelf	IOM	Port	Port pairs					
A and C	1	First	A	1	1a	2a	1c	2c		
	2	First	B	1						
B and D					~~1b~~	~~2b~~	~~1d~~	~~2d~~		
	1	Last	B	3	2b	1d	2d	1b		
	2	Last	A	3						

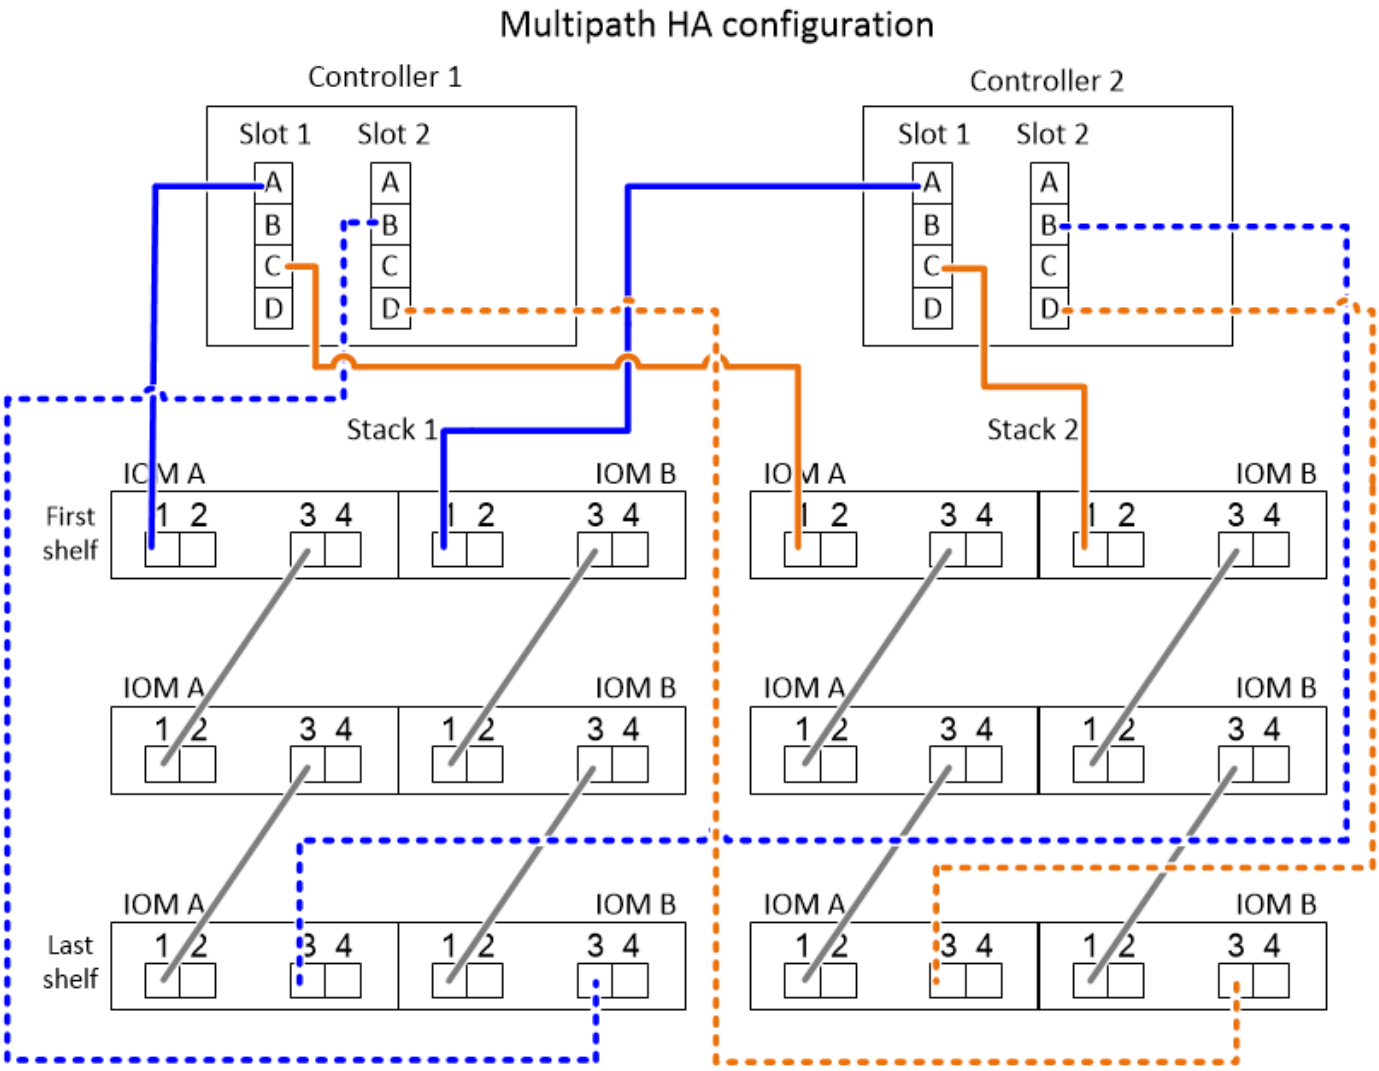

Steps

1. Cable port pair 1a/2b on each controller to stack 1:
   a. Cable controller 1 port 1a to stack 1, first shelf IOM A port 1.
   b. Cable controller 2 port 1a to stack 1, first shelf IOM B port 1.
   c. Cable controller 1 port 2b to stack 1, last shelf IOM B port 3.
   d. Cable controller 2 port 2b to stack 1, last shelf IOM A port 3.
2. Cable port pair 1c/2d on each controller to stack 2:
   a. Cable controller 1 port 1c to stack 2, first shelf IOM A port 1.
   b. Cable controller 2 port 1c to stack 2, first shelf IOM B port 1.
   c. Cable controller 1 port 2d to stack 2, last shelf IOM B port 3.
   d. Cable controller 2 port 2d to stack 2, last shelf IOM A port 3.

How to read a worksheet to cable controller-to-stack connections for quad-pathed connectivity - DS212C, DS224C, or DS460C

You can use this example to guide you through how to read and apply a completed worksheet to cable stacks of disk shelves with IOM12/IOM12B modules for quad-pathed connectivity.

About this task

- This procedure references the following worksheet and cabling example to demonstrate how to read a worksheet to cable controller-to-stack connections.

  The configuration used in this example is a quad-path HA configuration with two quad-port SAS HBAs on each controller and two stacks of disk shelves with IOM12 modules.

- If you have a single-controller configuration, skip substeps b and d for cabling to a second controller.
- If needed, you can refer to SAS cabling rules and concepts for information about the controller slot numbering convention, shelf-to-shelf connectivity, and controller-to-shelf connectivity (including the use of port pairs).

Controller-to-Stack Cabling Worksheet for Quad-Pathed Connectivity									
Controller SAS ports	Controllers	Cable to disk shelf IOMs				Stacks			
		Shelf	IOM	Port		1		2	
				Multipathed	Quad-pathed	Port pairs			
A and C	1	First	A	1	2	1a	2a	1c	2c
	2	First	B	1	2				
B and D						~~1b~~	~~2b~~	~~1d~~	~~2d~~
	1	Last	B	3	4	2b	1d	2d	1b
	2	Last	A	3	4				

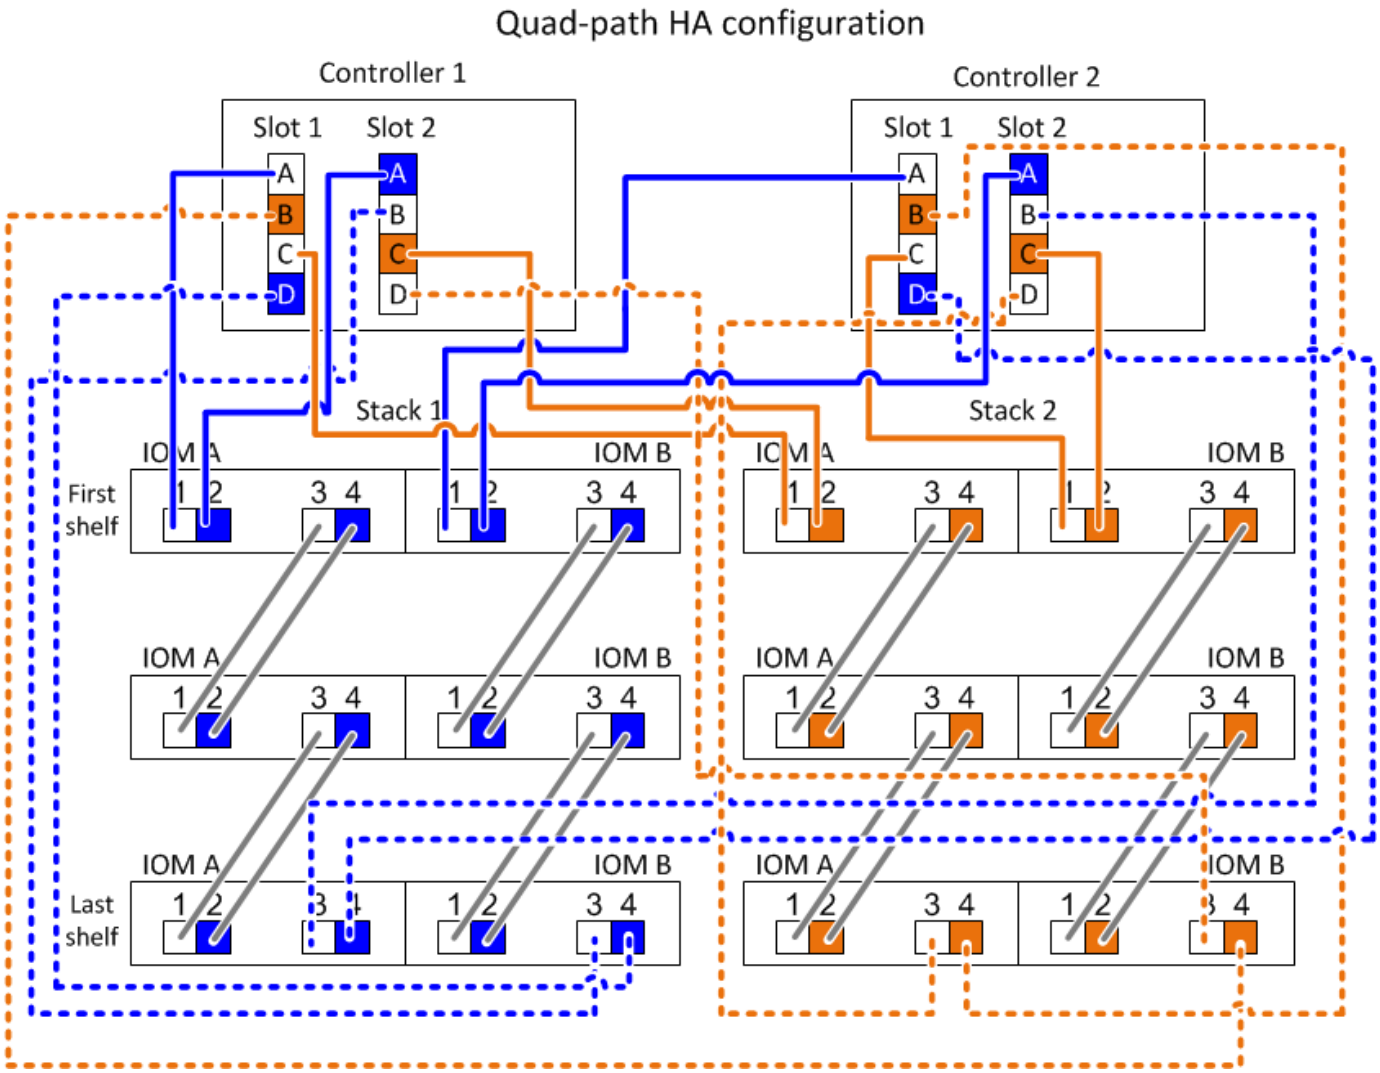

Steps

1. Cable port pair 1a/2b on each controller to stack 1:

   This is the multipathed cabling for stack 1.

   a. Cable controller 1 port 1a to stack 1, first shelf IOM A port 1.
   b. Cable controller 2 port 1a to stack 1, first shelf IOM B port 1.
   c. Cable controller 1 port 2b to stack 1, last shelf IOM B port 3.
   d. Cable controller 2 port 2b to stack 1, last shelf IOM A port 3.

2. Cable port pair 2a/1d on each controller to stack 1:

   This is the quad-pathed cabling for stack 1. Once completed, stack 1 has quad-pathed connectivity to each controller.

   a. Cable controller 1 port 2a to stack 1, first shelf IOM A port 2.
   b. Cable controller 2 port 2a to stack 1, first shelf IOM B port 2.
   c. Cable controller 1 port 1d to stack 1, last shelf IOM B port 4.
   d. Cable controller 2 port 1d to stack 1, last shelf IOM A port 4.

3. Cable port pair 1c/2d on each controller to stack 2:

   This is the multipathed cabling for stack 2.

   a. Cable controller 1 port 1c to stack 2, first shelf IOM A port 1.
   b. Cable controller 2 port 1c to stack 2, first shelf IOM B port 1.
   c. Cable controller 1 port 2d to stack 2, last shelf IOM B port 3.
   d. Cable controller 2 port 2d to stack 2, last shelf IOM A port 3.

4. Cable port pair 2c/1b on each controller to stack 2:

   This is the quad-pathed cabling for stack 2. Once completed, stack 2 has quad-pathed connectivity to each controller.

   a. Cable controller 1 port 2c to stack 2, first shelf IOM A port 2.
   b. Cable controller 2 port 2c to stack 2, first shelf IOM B port 2.
   c. Cable controller 1 port 1b to stack 2, last shelf IOM B port 4.
   d. Cable controller 2 port 1b to stack 2, last shelf IOM A port 4.

Maintain

Hot-swap a disk drive - DS212C, DS224C

You can hot-swap a failed disk drive in a DS224C or DS212C disk shelf with IOM12, IOM12B modules.

About this task

- Disk drive firmware is automatically updated (nondisruptively) on new disk drives with non current firmware versions.

  Disk drive firmware checks occur every two minutes.

- If needed, you can turn on the disk shelf's location (blue) LEDs to aid in physically locating the affected disk shelf: `storage shelf location-led modify -shelf-name shelf_name -led-status on`

  A disk shelf has three location LEDs: one on the operator display panel and one on each shelf IOM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the off option.

- If needed, you can refer to the Monitoring disk shelf LEDs section for information about the meaning and location of disk shelf LEDs on the operator display panel and FRU components.

Before you begin

- The disk drive that you are installing must be supported by the DS224C or DS212C disk shelf. You can view supported disk drives for your configuration in the NetApp Hardware Universe.
- All other components in the system must be functioning properly; if not, contact technical support.
- The disk drive you are removing must be failed.

  You can verify the disk drive is failed by running the `storage disk show -broken` command. The failed disk drive appears in the list of failed disk drives. If it does not, you should wait, and run the

command again.

> Depending on the disk drive type and capacity, it can take up to several hours for the disk drive to appear in the list of failed disk drives.

- If you are replacing a self-encrypting disk (SED), you must follow the instructions for Replacing an SED in the ONTAP documentation for your version of ONTAP.

  Instructions in the NetApp encryption overview with the CLI documentation describe additional steps you must perform before and after replacing an SED.

- **Best practice**: Ensure your system can recognize and utilize newly qualified disk drives by downloading the current version of the Disk Qualification Package (DQP) before hot-swapping a drive.

  This allows you to avoid system event messages about having non-current disk drive information. You also avoid the possible prevention of disk partitioning because disk drives are not recognized. The DQP notifies you of non-current disk drive firmware.

- **Best practice**: Ensure your system has the current versions of disk shelf (IOM) firmware and disk drive firmware on your system before adding new disk shelves, shelf FRU components, or SAS cables. You can visit the NetApp Support Site to download disk shelf firmware and download disk drive firmware.
- You should take steps to avoid electrostatic discharge (ESD):
  - Keep the disk drive in the ESD bag until you are ready to install it.
  - Open the ESD bag by hand or cut the top off with a pair of scissors.

    > Do not insert a metal tool or knife into the ESD bag.

  - Always wear an ESD wrist strap grounded to an unpainted surface on your storage enclosure chassis.

    If a wrist strap is unavailable, touch an unpainted surface on your storage enclosure chassis before handling the disk drive.

- You should take steps to handle disk drives carefully:
  - Always use two hands when removing, installing, or carrying a disk drive to support its weight.

    > Do not place hands on the disk drive boards exposed on the underside of the disk drive carrier.

  - You should place disk drives on cushioned surfaces, and never stack disk drives on top of each other.
  - You should be careful not to bump disk drives against other surfaces.
- Disk drives should be kept away from magnetic devices.

> Magnetic fields can destroy all data on the disk drive and cause irreparable damage to the disk drive circuitry.

Steps

1. If you want to manually assign disk ownership for the replacement disk drive, you need to disable automatic drive assignment if it is enabled; otherwise, go to the next step.

You need to manually assign disk ownership if disk drives in the stack are owned by both controllers in an HA pair.

You manually assign disk ownership and then reenable automatic drive assignment later in this procedure.

a. Verify if automatic drive assignment is enabled:`storage disk option show`

   If you have an HA pair, you can enter the command at the console of either controller.

   If automatic drive assignment is enabled, the output shows “on” (for each controller) in the “Auto Assign” column.

b. If automatic drive assignment is enabled, you need to disable it:`storage disk option modify -node` *`node_name`* `-autoassign off`

   You need to disable automatic drive assignment on both controllers in an HA pair.

2. Properly ground yourself.
3. Unpack the new disk drive, and set it on a level surface near the disk shelf.

   Save all packaging materials for use when returning the failed disk drive.

   NetApp requires that all returned disk drives be in a ESD-rated bag.

4. Physically identify the failed disk drive from the system console warning message and the illuminated attention (amber) LED on the disk drive.

   The activity (green) LED on a failed disk drive can be illuminated (solid), which indicates the disk drive has power, but should not be blinking, which indicates I/O activity. A failed disk drive has no I/O activity.

5. Press the release button on the disk drive face, and then pull the cam handle to its fully open position to release the disk drive from the mid plane.

   When you press the release button, the cam handle on the disk drive springs open partially.

   Disk drives in a DS212C disk shelf are arranged horizontally with the release button located on the left of the disk drive face. Disk drives in a DS224C disk shelf are arranged vertically with the release button located at the top of the disk drive face.

   The following shows disk drives in a DS212C disk shelf:

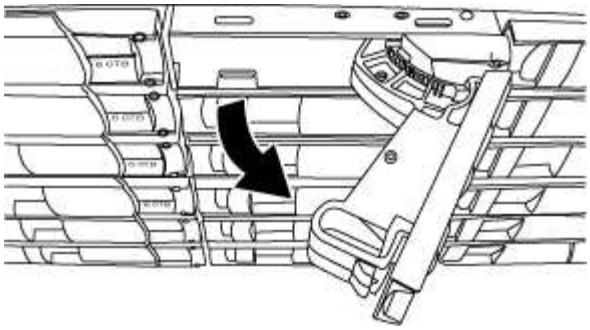

The following shows disk drives in a DS224C disk shelf:

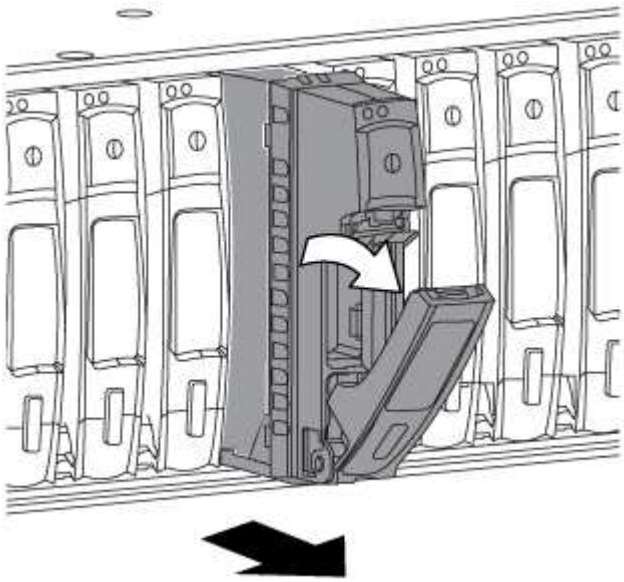

6. Slide out the disk drive slightly to allow the disk to safely spin down, and then remove the disk drive from the disk shelf.

   An HDD can take up to one minute to safely spin down.

   When handling a disk drive, always use two hands to support its weight.

7. Using two hands, with the cam handle in the open position, insert the replacement disk drive into the disk shelf, firmly pushing until the disk drive stops.

   Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   Do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

8. Close the cam handle so that the disk drive is fully seated into the mid plane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive.

9. If you are replacing another disk drive, repeat Steps 3 through 8.

10. Verify the disk drive's activity (green) LED is illuminated.

    When the disk drive's activity LED is solid green, it means the disk drive has power. When the disk drive's activity LED is blinking, it means the disk drive has power and I/O is in progress. If the disk drive firmware is automatically updating, the LED will be blinking.

11. If you disabled automatic drive assignment in Step 1, manually assign disk ownership, and then reenable automatic drive assignment if needed:

    a. Display all unowned disks:`storage disk show -container-type unassigned`

    b. Assign each disk:`storage disk assign -disk disk_name -owner owner_name`

       You can use the wildcard character to assign more than one disk at once.

    c. Reenable automatic drive assignment if needed:`storage disk option modify -node node_name -autoassign on`

       You need to reenable automatic drive assignment on both controllers in an HA pair.

12. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Hot-swap a disk drive - DS460C

You can hot-swap a failed disk drive in a DS460C disk shelf with IOM12 or IOM12B modules.

Before you begin

- The replacement disk drive must be supported by the DS460C disk shelf. You can view supported disk drives for your configuration in the NetApp Hardware Universe.
- All other components in the system must be functioning properly; if not, contact technical support.
- The disk drive you are removing must be failed.

  You can verify the disk drive is failed by running the `storage disk show -broken` command. The failed disk drive appears in the list of failed disk drives. If it does not, you should wait, and run the command again.

  > Depending on the disk drive type and capacity, it can take up to several hours for the disk drive to appear in the list of failed disk drives.

- If you are replacing a self-encrypting disk (SED), you must follow the instructions for Replacing an SED in the ONTAP documentation for your version of ONTAP.

  Instructions in the NetApp encryption overview with the CLI documentation describe additional steps you must perform before and after replacing an SED.

About this task

- You should take steps to avoid electrostatic discharge (ESD):

- Keep the disk drive in the ESD bag until you are ready to install it.
- Open the ESD bag by hand or cut the top off with a pair of scissors.

  Do not insert a metal tool or knife into the ESD bag.

- Always wear an ESD wrist strap grounded to an unpainted surface on your storage enclosure chassis.

  If a wrist strap is unavailable, touch an unpainted surface on your storage enclosure chassis before handling the disk drive.

- You should take steps to handle disk drives carefully:
  - Always use two hands when removing, installing, or carrying a disk drive to support its weight.

    Do not place hands on the disk drive boards exposed on the underside of the disk drive carrier.

  - You should place disk drives on cushioned surfaces, and never stack disk drives on top of each other.
  - You should be careful not to bump disk drives against other surfaces.
- Disk drives should be kept away from magnetic devices.

  Magnetic fields can destroy all data on the disk drive and cause irreparable damage to the disk drive circuitry.

- **Best practice**: Ensure your system can recognize and utilize newly qualified disk drives by downloading the current version of the Disk Qualification Package (DQP) before hot-swapping a disk drive.

  Having the current version of the DQP installed allows your system to recognize and utilize newly qualified disk drives; therefore, avoiding system event messages about having non-current disk drive information. You also avoid the possible prevention of disk partitioning because disk drives are not recognized. The DQP notifies you of non-current disk drive firmware.

- **Best practice**: Ensure your system has the current versions of disk shelf (IOM) firmware and disk drive firmware on your system before adding new disk shelves, shelf FRU components, or SAS cables. You can visit the NetApp Support Site to download disk shelf firmware and download disk drive firmware.
- Disk drive firmware is automatically updated (nondisruptively) on new disk drives with non current firmware versions.

  Disk drive firmware checks occur every two minutes.

- If needed, you can turn on the disk shelf's location (blue) LEDs to aid in physically locating the affected disk shelf: `storage shelf location-led modify -shelf-name shelf_name -led-status on`

  A disk shelf has three location LEDs: one on the operator display panel and one on each shelf IOM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the off option.

- If needed, you can refer to the Monitoring disk shelf LEDs section for information about the meaning and location of disk shelf LEDs on the operator display panel and FRU components.
- The DS460C drive shelf consist of five drive drawers (drive drawer 1 at the top through drive drawer 5 at the bottom) that each contain 12 drive slots.

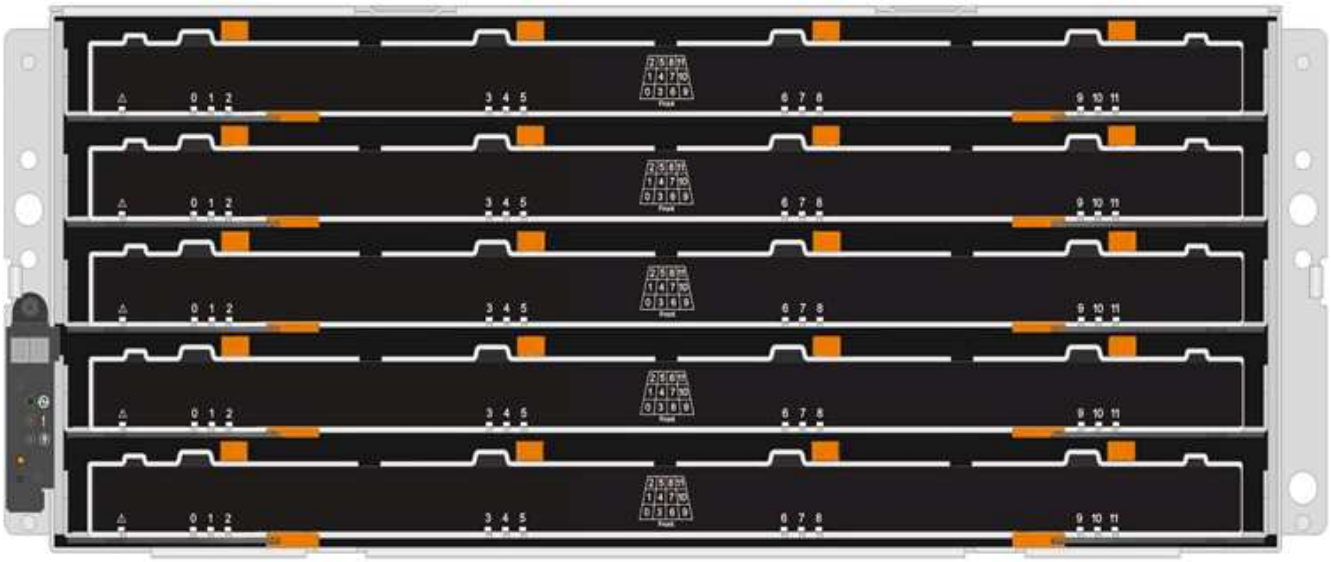

- The following illustration shows how the drives are numbered from 0 to 11 in each drive drawer within the shelf.

Steps

1. If you want to manually assign disk ownership for the replacement disk drive, you need to disable automatic drive assignment if it is enabled; otherwise, go to the next step.

   You need to manually assign disk ownership if disk drives in the stack are owned by both controllers in an HA pair.

   You manually assign disk ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify if automatic drive assignment is enabled:`storage disk option show`

      If you have an HA pair, you can enter the command at the console of either controller.

If automatic drive assignment is enabled, the output shows “on” (for each controller) in the “Auto Assign” column.

b. If automatic drive assignment is enabled, you need to disable it:`storage disk option modify -node` *`node_name`* `-autoassign off`

You need to disable automatic drive assignment on both controllers in an HA pair.

2. Properly ground yourself.
3. Unpack the new disk drive, and set it on a level surface near the disk shelf.

Save all packaging materials for use when returning the failed disk drive.

NetApp requires that all returned disk drives be in a ESD-rated bag.

4. Identify the failed disk drive from the system console warning message and the illuminated amber attention LED on the drive drawer.

The 2.5-inch and 3.5-inch SAS drive carriers do not contain LEDs. Instead, you must look at the Attention LEDs on the drive drawers to determine which drive has failed.

The drive drawer’s Attention LED (amber) blinks so you can open the correct drive drawer to identify which drive to replace.

The drive drawer’s Attention LED is on the front-left side in front of each drive, with a warning symbol on the drive handle just behind the LED.

5. Open the drawer containing the failed drive:
   a. Unlatch the drive drawer by pulling on both levers.
   b. Using the extended levers, carefully pull the drive drawer out until it stops.
   c. Look at the top of the drive drawer to find the Attention LED that resides on the drawer in front of each drive.
6. Remove the failed drive from the open drawer:
   a. Gently pull back the orange release latch that is in front of the drive you want to remove.

1	Orange release latch

b. Open the cam handle, and lift out the drive slightly.
c. Wait 30 seconds.
d. Use the cam handle to lift the drive from the shelf.

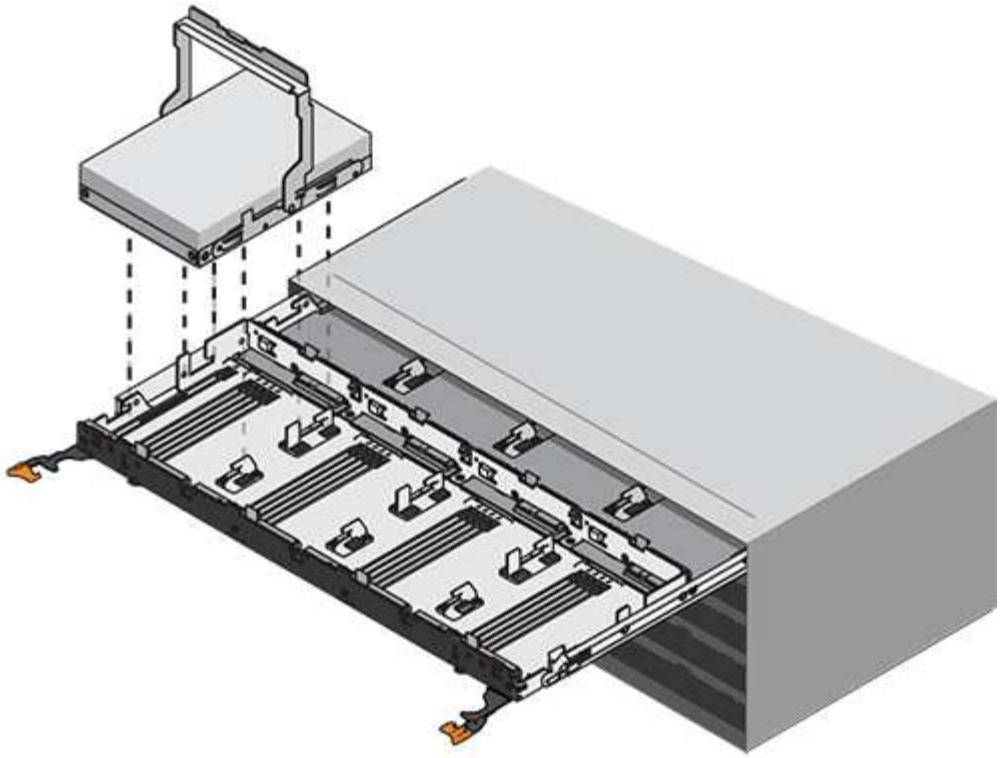

e. Place the drive on an antistatic, cushioned surface away from magnetic fields.

7. Insert the replacement drive in the drawer:
   a. Raise the cam handle on the new drive to vertical.
   b. Align the two raised buttons on each side of the drive carrier with the matching gap in the drive channel on the drive drawer.

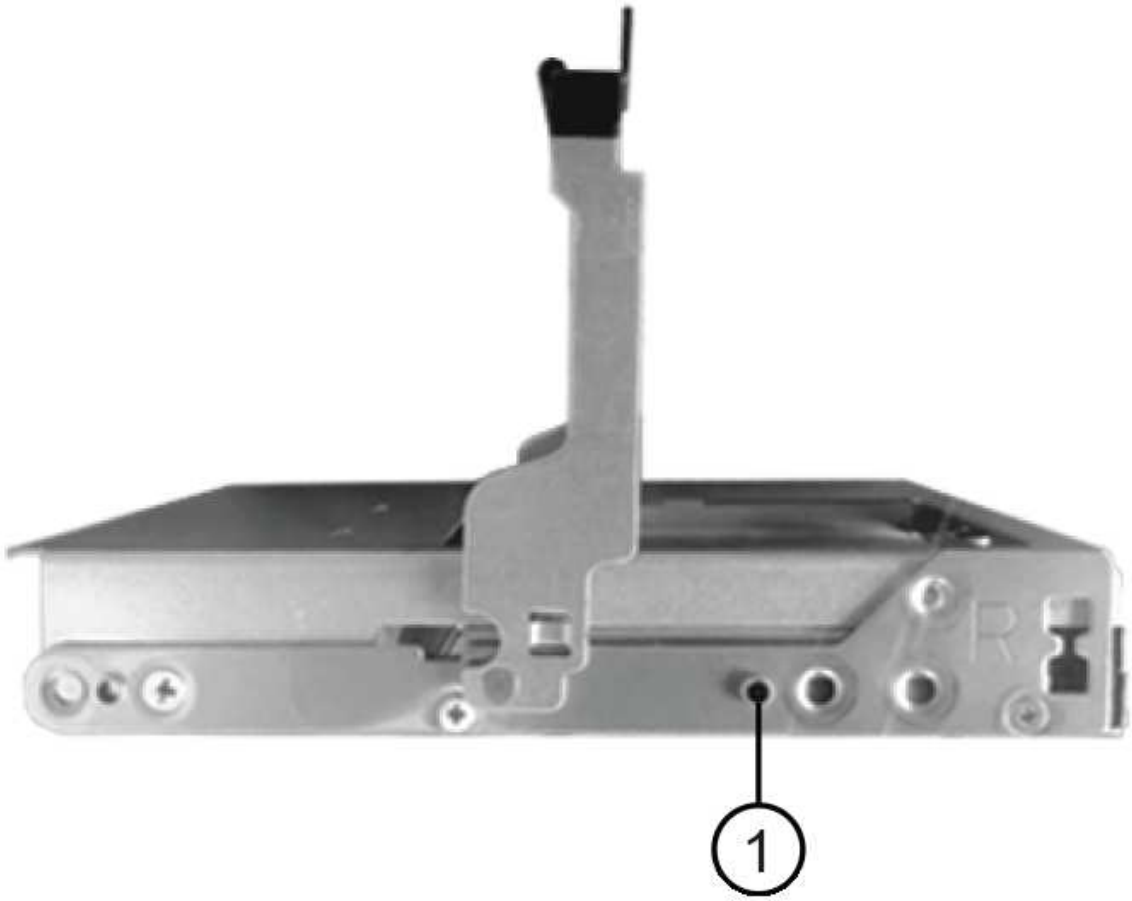

1	Raised button on the right side of the drive carrier

c. Lower the drive straight down, and then rotate the cam handle down until the drive snaps into place under the orange release latch.

d. Carefully push the drive drawer back into the enclosure.

Possible loss of data access: Never slam the drawer shut. Push the drawer in slowly to avoid jarring the drawer and causing damage to the storage array.

e. Close the drive drawer by pushing both levers towards the center.

   The green Activity LED for the replaced drive on the front of the drive drawer comes on when the drive is inserted correctly.

8. If you are replacing another disk drive, repeat Steps 4 through 7.
9. Check the Activity LED and the Attention LED on the drive you replaced.

LED status	Description
The Activity LED is on or blinking, and the Attention LED is off	The new drive is working correctly.
The Activity LED is off	The drive might not be installed correctly. Remove the drive, wait 30 seconds, and then reinstall it.
The Attention LED is on	The new drive might be defective. Replace it with another new drive. When you first insert a drive, its Attention LED might be on. However, the LED should go off within a minute.

10. If you disabled disk ownership automatic assignment in Step 1, manually assign disk ownership, and then reenable disk ownership automatic assignment if needed:
    a. Display all unowned disks:`storage disk show -container-type unassigned`
    b. Assign each disk:`storage disk assign -disk disk_name -owner owner_name`

       You can use the wildcard character to assign more than one disk at once.

    c. Reenable disk ownership automatic assignment if needed:`storage disk option modify -node node_name -autoassign on`

       You need to reenable disk ownership automatic assignment on both controllers in an HA pair.

11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a drive drawer - DS460C

To replace a drive drawer in a shelf, you can halt the entire storage system (HA pair), which allows you to keep the data aggregates on the shelf. Alternatively, you can keep the HA pair up and running, which requires you to move all data from the data aggregates on the shelf's disk drives and then offline and delete those data aggregates. If the shelf contains a root aggregate, you must halt the HA pair.

This procedure applies to shelves with IOM12 or IOM12B modules.

Before you begin

You need these items for this procedure:

- Antistatic protection

> **Possible hardware damage:** To prevent electrostatic discharge damage to the drive shelf, use proper antistatic protection when handling drive shelf components.

- Replacement drive drawer
- Replacement left and right cable chains
- Flashlight

About this task

- This procedure applies to shelves having DCM drive drawers and/or DCM2 or DCM3 drive drawers. (Shelves will also have two IOM12 modules or two IOM12B modules.)

  When a DCM, DCM2, or DCM3 drive drawer fails, you receive a DCM, DCM2, or DCM3 drive drawer to replace it.

> When you replace a failed drive drawer with a newer DCM2 or DCM3 drive drawer, ensure that ONTAP software and the IOM12 or IOM12B module firmware are running the minimum versions required to support DCM2 or DCM3 drawers.
>
> The IOM12 FW upgrade can be done before or after replacing a drive drawer. This procedure has you upgrade the FW as part of the preparation for the drawer replacement procedure.

- The DCM, DCM2, and DCM3 drive drawers can be distinguished by their appearance:

  The DCM drive drawers look like the following:

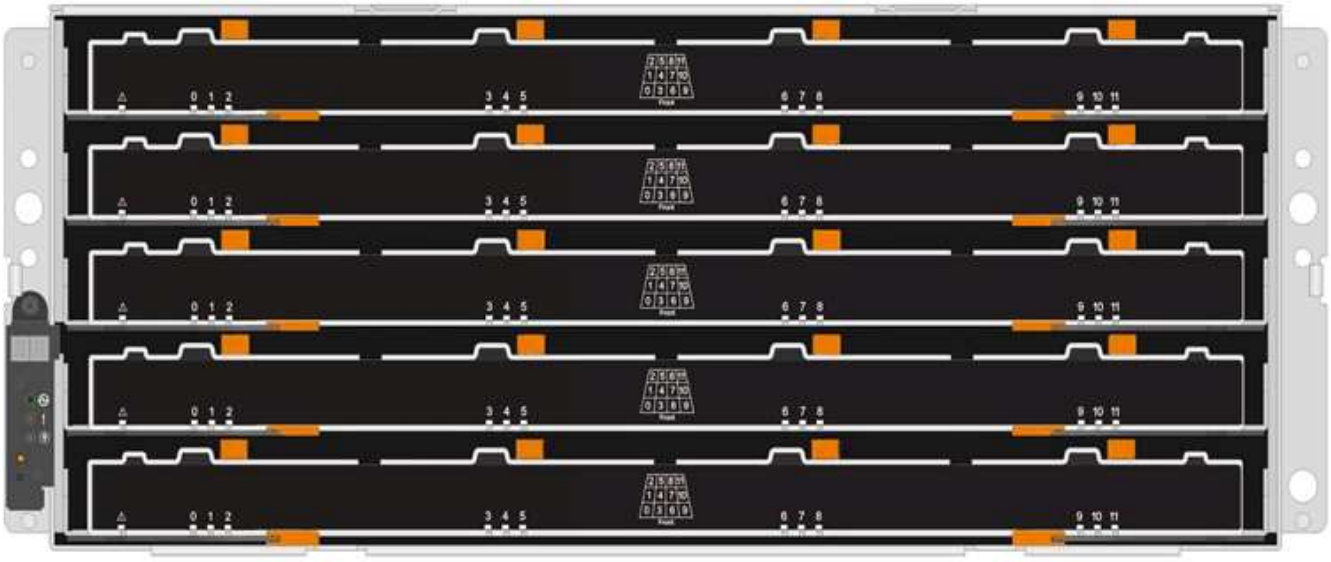

  The DCM2 drive drawers are distinguished by a blue stripe and "DCM2" label:

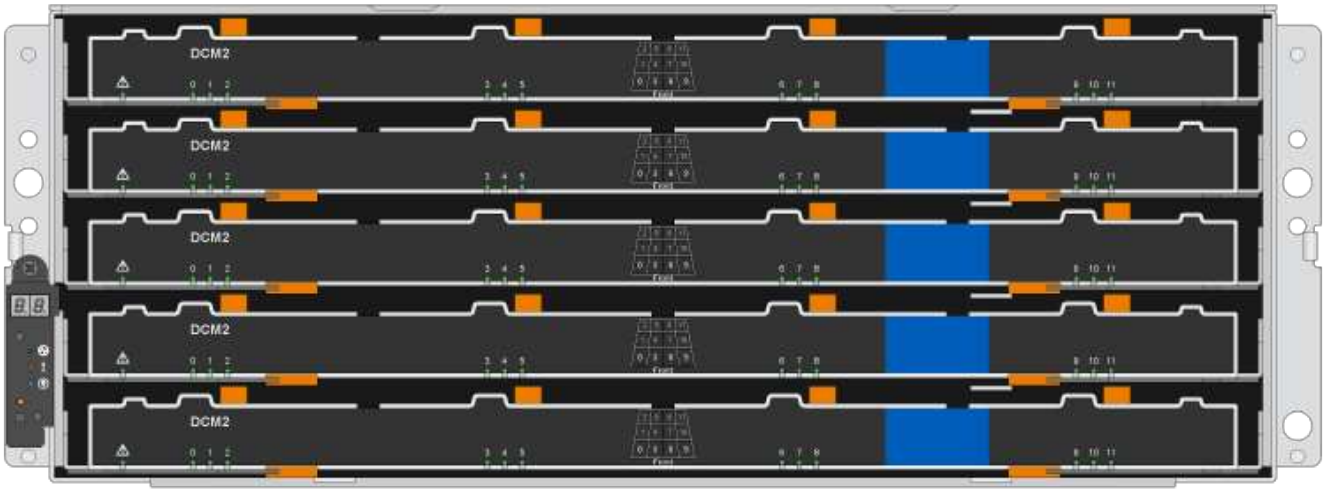

The DCM3 drive drawers are distinguished by a blue and gray stripe and "DCM3" label:

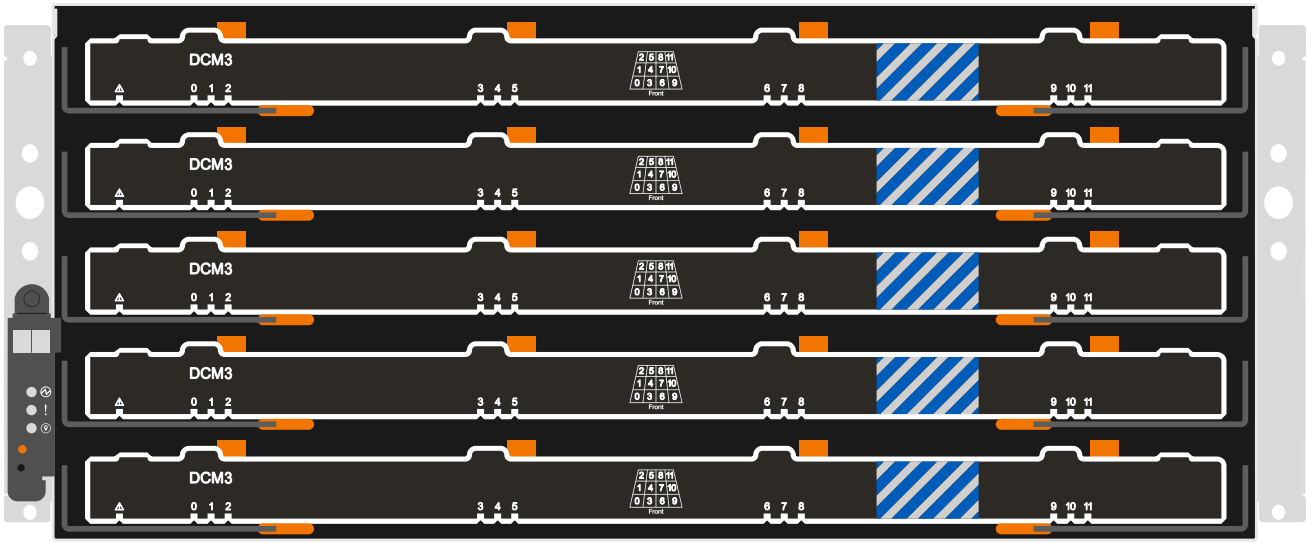

Step 1: Prepare to replace a drive drawer

Before you replace a drive drawer, you must make sure to update IOM FW and ONTAP if needed, and either halt the HA pair, which allows you to keep data aggregates on the shelf, or keep the HA pair up and running, which requires you to move all of the data from the data aggregates residing on the disk drives. You must then offline and delete the data aggregates. However, if the shelf contains a root aggregate, you must halt the HA pair. Lastly, you must power off the shelf.

Steps

1. Determine if when you replace the failed drive drawer, it results in the shelf having a combination of IOM12 modules and any number of DCM2 drive drawers.
2. Determine if you will need to upgrade your IOM firmware:
   - If the shelf will have a combination of IOM12 modules and any number of DCM2 drive drawers, you must upgrade the IOM12 FW to version 0300 or later; otherwise, go to the next step.
   - If the shelf will have a combination of IOM12 modules and any number of DCM3 drive drawers, you must upgrade the IOM12 FW to version 0401 or later.

- If the shelf will have a combination of IOM12B modules and any number of DCM3 drive drawers, you must upgrade the IOM12B FW to version 0202 or later.

If needed, you can go to the NetApp Support Site to download current versions of disk shelf firmware. Otherwise, go to the next step.

3. If your shelf will have a combination of IOM12 or IOM12B modules with any number of DCM3 drive drawers, your system must be running the following minimum versions of ONTAP:
   - ONTAP 9.17.1RC1
   - ONTAP 9.16.1P3
   - ONTAP 9.15.1P11
   - ONTAP 9.14.1P13
   - ONTAP 9.13.1P16
   - ONTAP 9.12.1P18
4. If the shelf contains the root aggregate from either controller in the HA pair, or if you chose to halt the HA pair (instead of keeping the HA pair up and running), complete the following substeps; otherwise, go to the next step.

   Data aggregates can remain on the shelf when halting the HA pair.

   a. Halt both controllers in the HA pair.
   b. Verify that your HA pair halted by checking the storage system console.
   c. Power off the shelf.
   d. Go to the section, Remove the cable chains.
5. If you chose to keep the HA pair up and running, complete the following substeps:

   If you attempt to replace a drawer with aggregates on the disk shelf, you could cause a system disruption with a multidisk panic.

   a. Move all of the data from the data aggregates located on all disk drives on the shelf.

      Data includes, but is not limited to, volumes and LUNs.

   b. Offine and delete the aggregates to allow the disk drives to become spares:

      Commands can be entered from the clustershell of either controller.

      `storage aggregate offline -aggregate` *`aggregate_name`*

      `storage aggregate delete -aggregate` *`aggregate_name`*

   c. Verify that the disk drives have no aggregates (are spares):
      i. Enter the following command from the clustershell of either controller: `storage disk show -shelf` *`shelf_number`*
      ii. Check the output to verify that the drisk drives are spares.

      Disk drives that are spares show `spare` in the `Container Type` column.

If you have failed disk drives in the shelf, `broken` appears in the `Container Type` column.

d. Power off the shelf.

Step 2: Remove the cable chains

Left and right cable chains for each drive drawer in the DS460C drive shelf allow the drawers to slide in and out. Before you can remove a drive drawer, you must remove both cable chains.

About this task

Each drive drawer has left and right cable chains. The metal ends on the cable chains slide into corresponding vertical and horizontal brackets inside the enclosure, as follows:

- The left and right vertical brackets connect the cable chain to the enclosure's midplane.
- The left and right horizontal brackets connect the cable chain to the individual drawer.

Before you begin

- You have completed the Prepare to replace a drive drawer steps so that your HA pair is halted or you have moved all of the data from the data aggregates residing on the disk drives, and offlined and deleted the data aggregates to allow the disk drives to become spares.
- You have powered off the shelf.
- You have obtained the following items:
  - Antistatic protection

    Possible hardware damage: To prevent electrostatic discharge damage to the shelf, use proper antistatic protection when handling shelf components.

  - Flashlight

Steps

1. Put on antistatic protection.
2. From the rear of the drive shelf, remove the right fan module, as follows:
   a. Press the orange tab to release the fan module handle.

      The figure shows the handle for the fan module extended and released from the orange tab on the left.

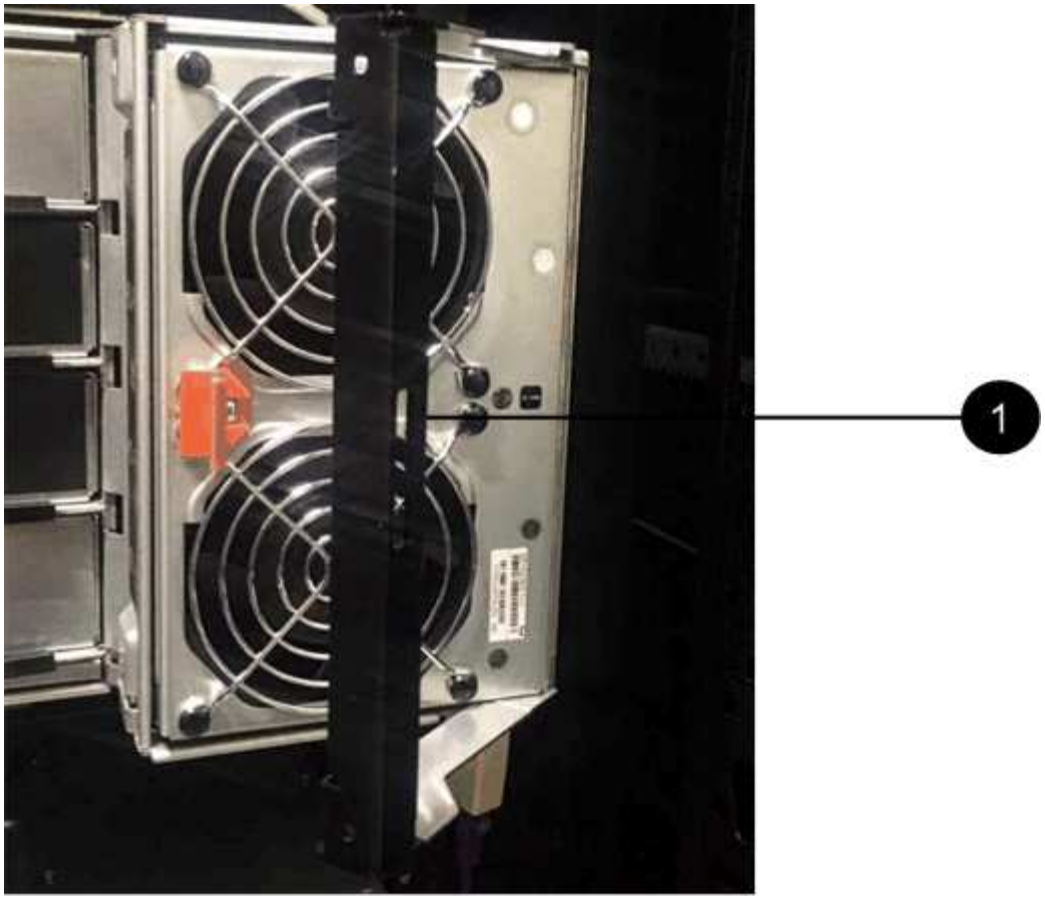

1	Fan module handle

b. Using the handle, pull the fan module out of the drive shelf, and set it aside.

3. Manually determine which of the five cable chains to disconnect.

   The figure shows the right side of the drive shelf with the fan module removed. With the fan module removed, you can see the five cable chains and the vertical and horizontal connectors for each drawer. The callouts for drive drawer 1 are provided.


2
1
3

1	Cable chain
2	Vertical connector (connected to the midplane)
3	Horizontal connector (connected to the drive drawer)

The top cable chain is attached to drive drawer 1. The bottom cable chain is attached to drive drawer 5.

4. Use your finger to move the cable chain on the right side to the left.
5. Follow these steps to disconnect any of the right cable chains from its corresponding vertical bracket.
   a. Using a flashlight, locate the orange ring on the end of the cable chain that is connected to the vertical bracket in the enclosure.

1	Orange ring on the vertical bracket

b. Disconnect the vertical connector (connected to the midplane) by gently pressing on the center of the

orange ring and pulling the left side of the cable out of the enclosure.

c. To unplug the cable chain, carefully pull your finger toward you approximately 1 inch (2.5 cm), but leave the cable chain connector within the vertical bracket.

6. Follow these steps to disconnect the other end of the cable chain:

a. Using a flashlight, locate the orange ring on the end of the cable chain that is attached to the horizontal bracket in the enclosure.

The figure shows the horizontal connector on the right and the cable chain disconnected and partially pulled out on the left side.

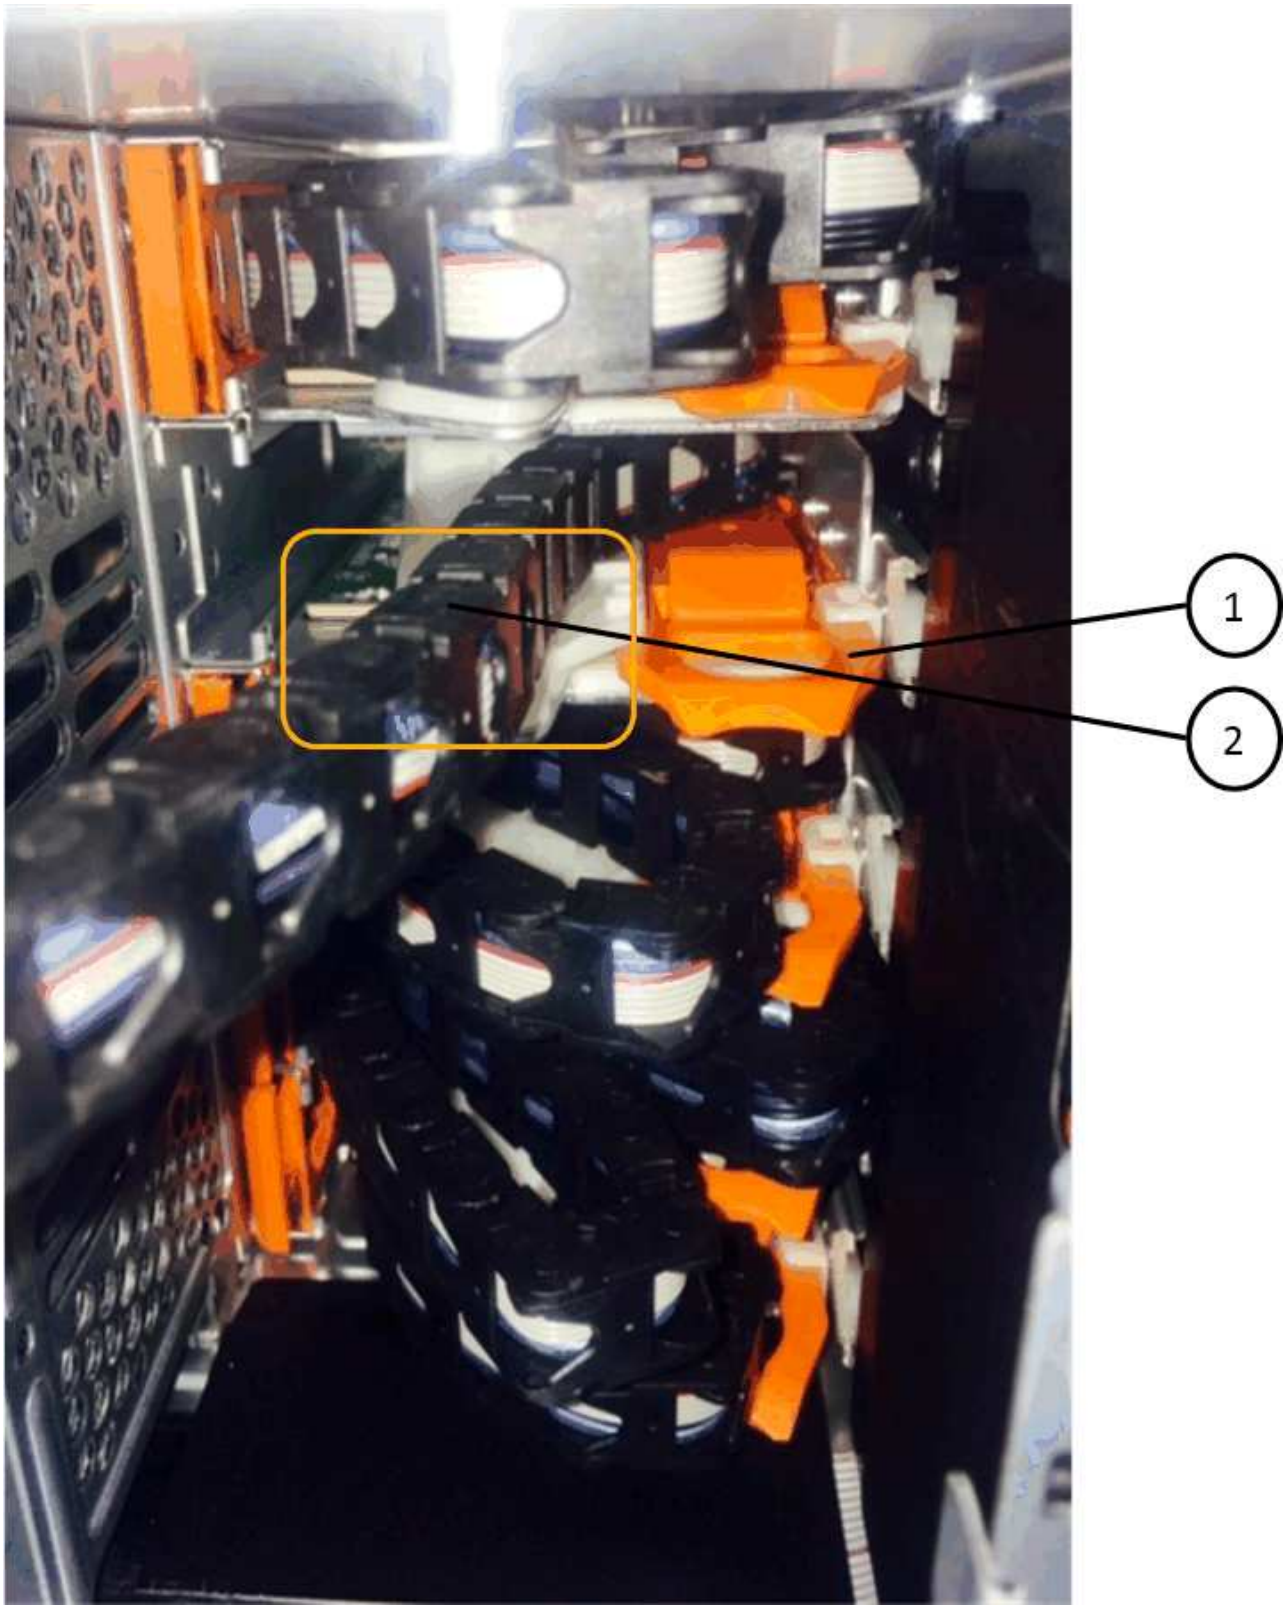

1	Orange ring on horizontal bracket
2	Cable chain

b. Gently insert your finger into the orange ring.

   The figure shows the orange ring on the horizontal bracket being pushed down so that the rest of the cable chain can be pulled out of the enclosure.

c. Pull your finger toward you to unplug the cable chain.

7. Carefully pull the entire cable chain out of the drive shelf.
8. From the back of the drive shelf, remove the left fan module.
9. Follow these steps to disconnect the left cable chain from its vertical bracket:
   a. Using a flashlight, locate the orange ring on the end of the cable chain attached to the vertical bracket.
   b. Insert your finger into the orange ring.
   c. To unplug the cable chain, pull your finger toward you approximately 1 inch (2.5 cm), but leave the cable chain connector within the vertical bracket.
10. Disconnect the left cable chain from the horizontal bracket, and pull the entire cable chain out of the drive shelf.

Step 3: Remove a drive drawer

After removing the right and left cable chains, you can remove the drive drawer from the drive shelf. Removing a drive drawer entails sliding the drawer part of the way out, removing the drives, and removing the drive drawer.

Before you begin

- You have removed the right and left cable chains for the drive drawer.
- You have replaced the right and left fan modules.

Steps

1. Remove the bezel from the front of the drive shelf.
2. Unlatch the drive drawer by pulling out on both levers.
3. Using the extended levers, carefully pull the drive drawer out until it stops. Do not completely remove the drive drawer from the drive shelf.
4. Remove the drives from the drive drawer:
   a. Gently pull back the orange release latch that is visible on the center front of each drive. The following image shows the orange release latch for each of the drives.

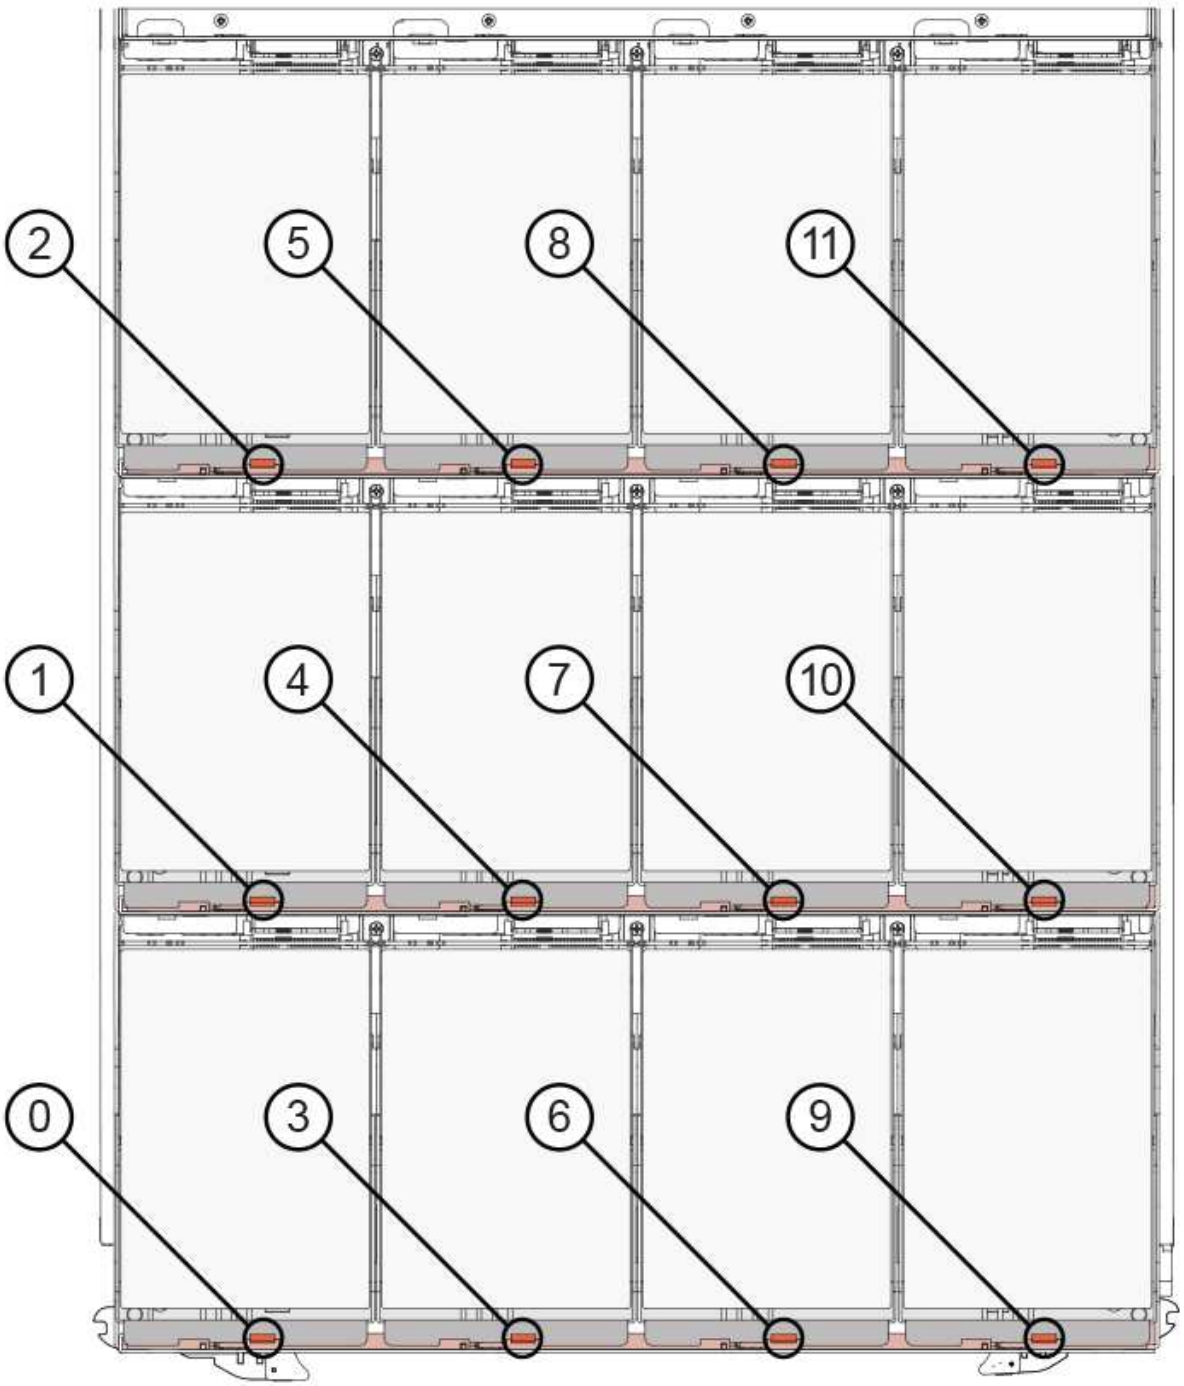

b. Raise the drive handle to vertical.

c. Use the handle to lift the drive from the drive drawer.

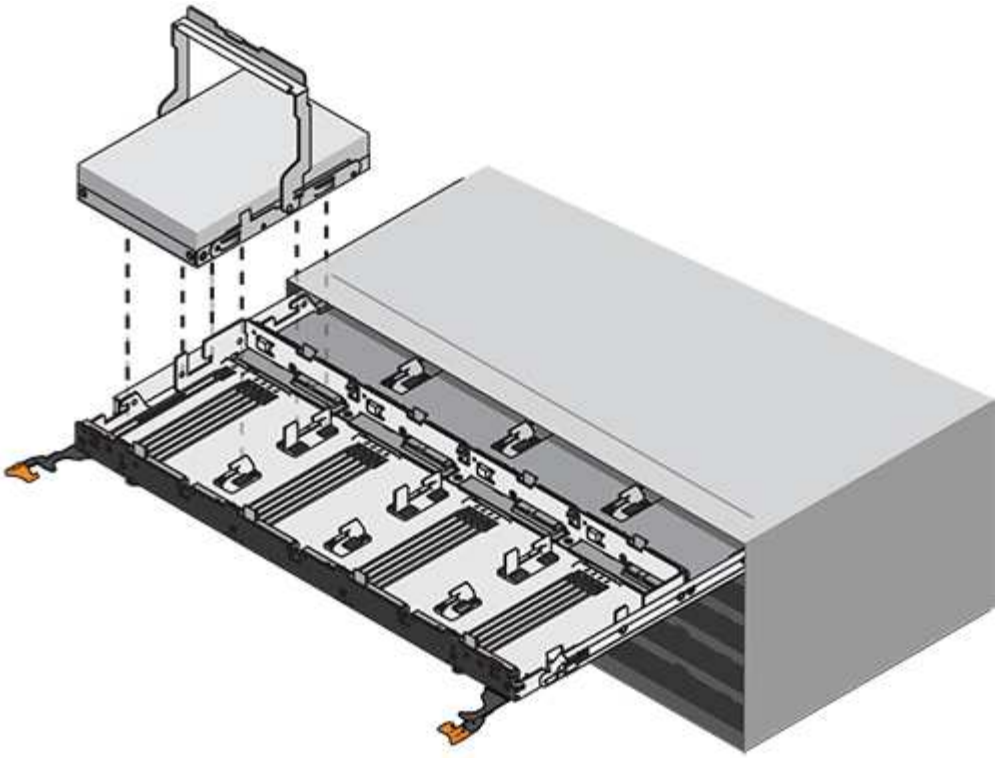

d. Place the drive on a flat, static-free surface and away from magnetic devices.

Possible loss of data access: Magnetic fields can destroy all data on the drive and cause irreparable damage to the drive circuitry. To avoid loss of data access and damage to the drives, always keep drives away from magnetic devices.

5. Follow these steps to remove the drive drawer:
   a. Locate the plastic release lever on each side of the drive drawer.

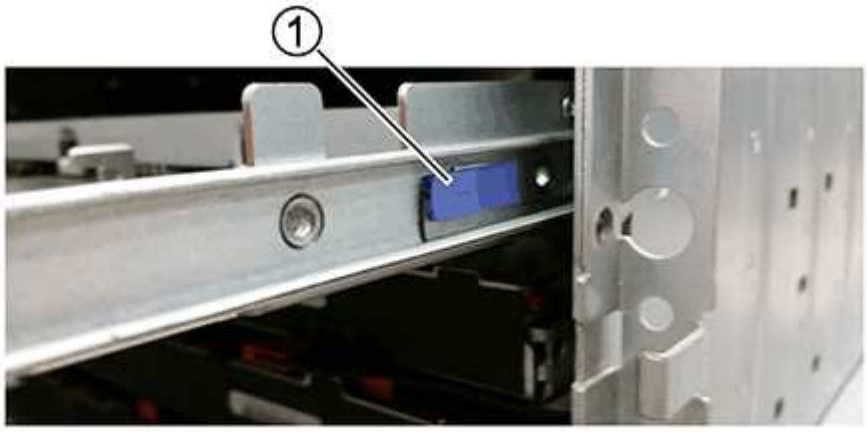

1	Drive drawer release lever

   b. Open both release levers by pulling the latches toward you.

c. While holding both release levers, pull the drive drawer toward you.

d. Remove the drive drawer from the drive shelf.

Step 4: Install a drive drawer

Installing a drive drawer into a drive shelf entails sliding the drawer into the empty slot, installing the drives, and replacing the front bezel.

Before you begin

- You have obtained the following items:
  - Replacement drive drawer
  - Flashlight

Steps

1. From the front of the drive shelf, shine a flashlight into the empty drawer slot, and locate the lock-out tumbler for that slot.

   The lock-out tumbler assembly is a safety feature that prevents you from being able to open more than one drive drawer at one time.

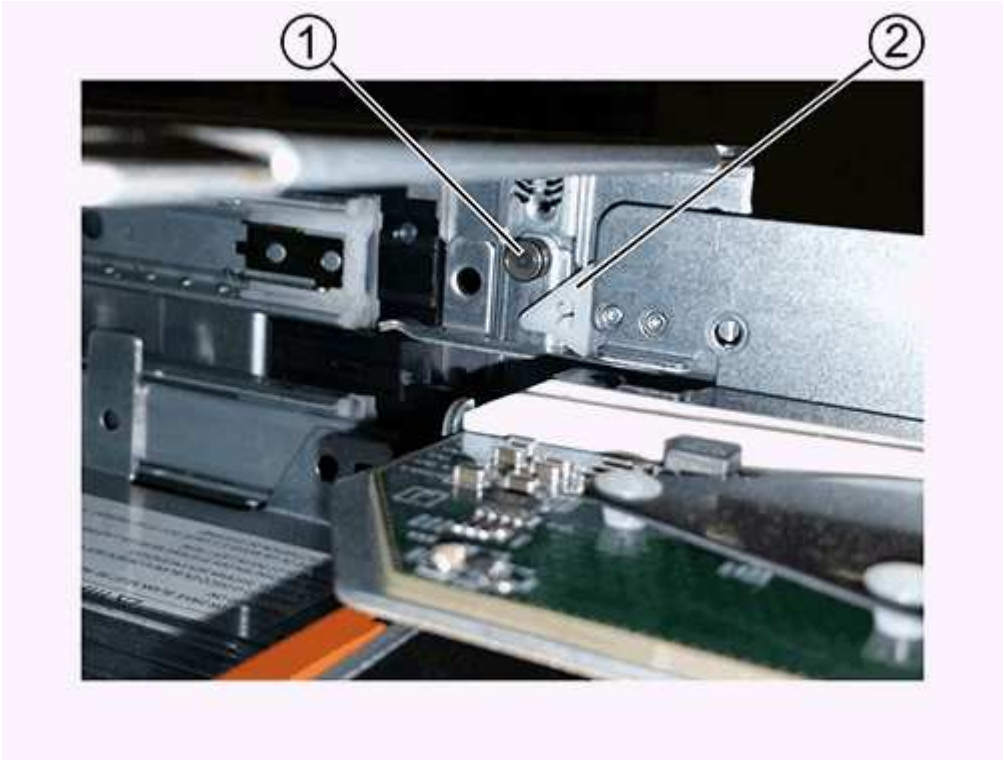

1	Lock-out tumbler
2	Drawer guide

2. Position the replacement drive drawer in front of the empty slot and slightly to the right of center.

   Positioning the drawer slightly to the right of center helps to ensure that the lock-out tumbler and the

drawer guide are correctly engaged.

3. Slide the drive drawer into the slot, and ensure that the drawer guide slides under the lock-out tumbler.

> **Risk of equipment damage:** Damage occurs if the drawer guide does not slide under the lock-out tumbler.

4. Carefully push the drive drawer all the way in until the latch fully engages.

> **Risk of equipment damage:** Stop pushing the drive drawer if you feel excessive resistance or binding. Use the release levers at the front of the drawer to slide the drawer back out. Then, reinsert the drawer into the slot, and ensure that it slides in and out freely.

5. Follow these steps to reinstall the drives in the drive drawer:
   a. Unlatch the drive drawer by pulling out on both levers at the front of the drawer.
   b. Using the extended levers, carefully pull the drive drawer out until it stops. Do not completely remove the drive drawer from the drive shelf.
   c. On the drive you are installing, raise the handle to vertical.
   d. Align the two raised buttons on each side of the drive with the notches on the drawer.

   The figure shows the right side view of a drive, showing the location of the raised buttons.

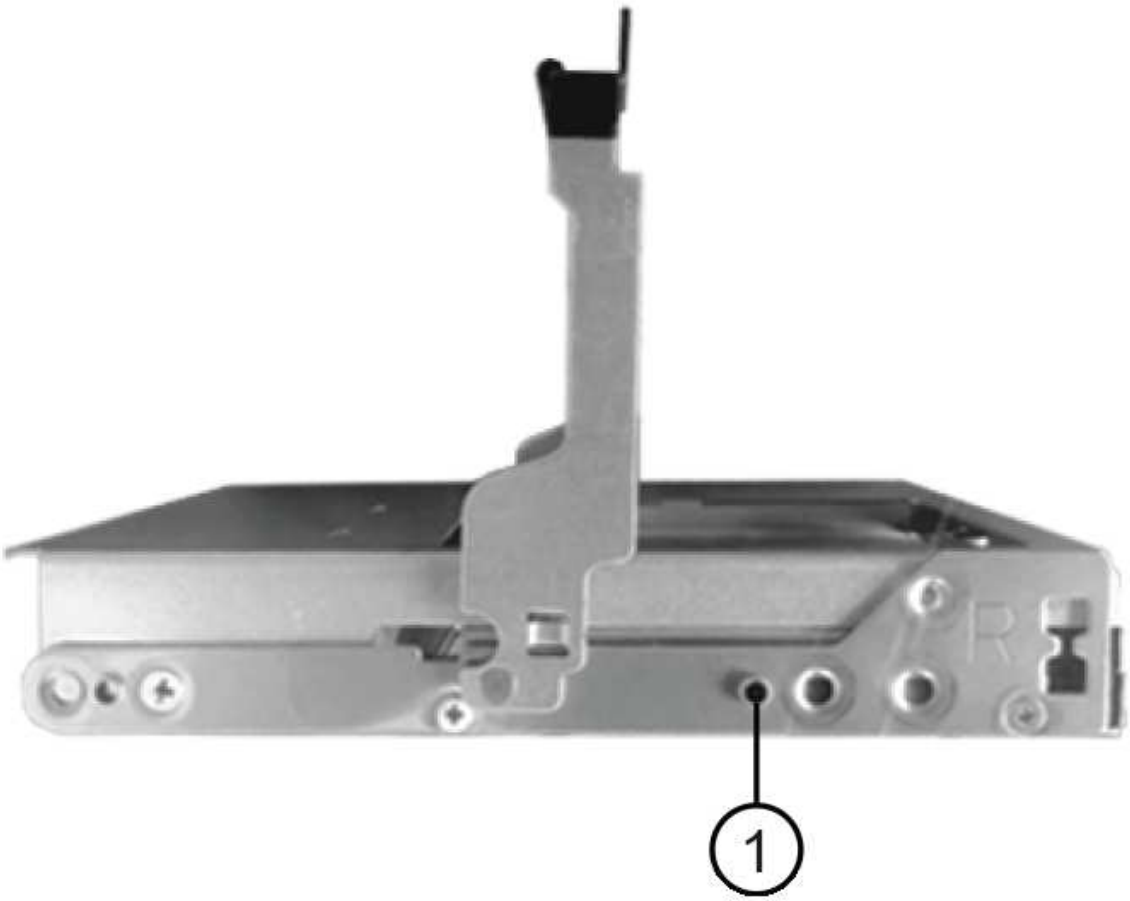

1	Raised button on the right side of the drive.

e. Lower the drive straight down, and then rotate the drive handle down until the drive snaps into place.

If you have a partially populated shelf, meaning that the drawer in which you are reinstalling drives has less than the 12 drives it supports, install the first four drives into the front slots (0, 3, 6, and 9).

Risk of equipment malfunction: To allow for proper air flow and prevent overheating, always install the first four drives into the front slots (0, 3, 6, and 9).

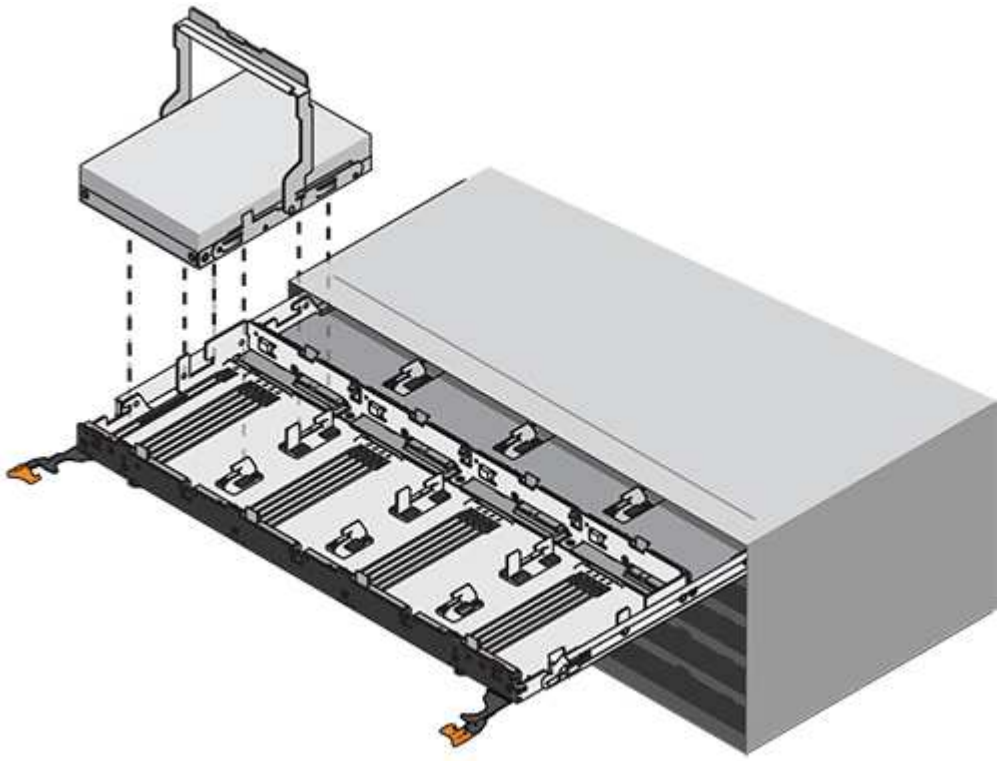

f. Repeat these substeps to reinstall all of the drives.

6. Slide the drawer back into the drive shelf by pushing it from the center and closing both levers.

Risk of equipment malfunction: Make sure to completely close the drive drawer by pushing both levers. You must completely close the drive drawer to allow proper airflow and prevent overheating.

7. Attach the bezel to the front of the drive shelf.

Step 5: Attach the cable chains

The final step in installing a drive drawer is attaching the replacement left and right cable chains to the drive shelf. When attaching a cable chain, reverse the order you used when disconnecting the cable chain. You must insert the chain's horizontal connector into the horizontal bracket in the enclosure before inserting the chain's vertical connector into the vertical bracket in the enclosure.

Before you begin

- You have replaced the drive drawer and all of the drives.

- You have two replacement cable chains, marked as LEFT and RIGHT (on the horizontal connector next to the drive drawer).

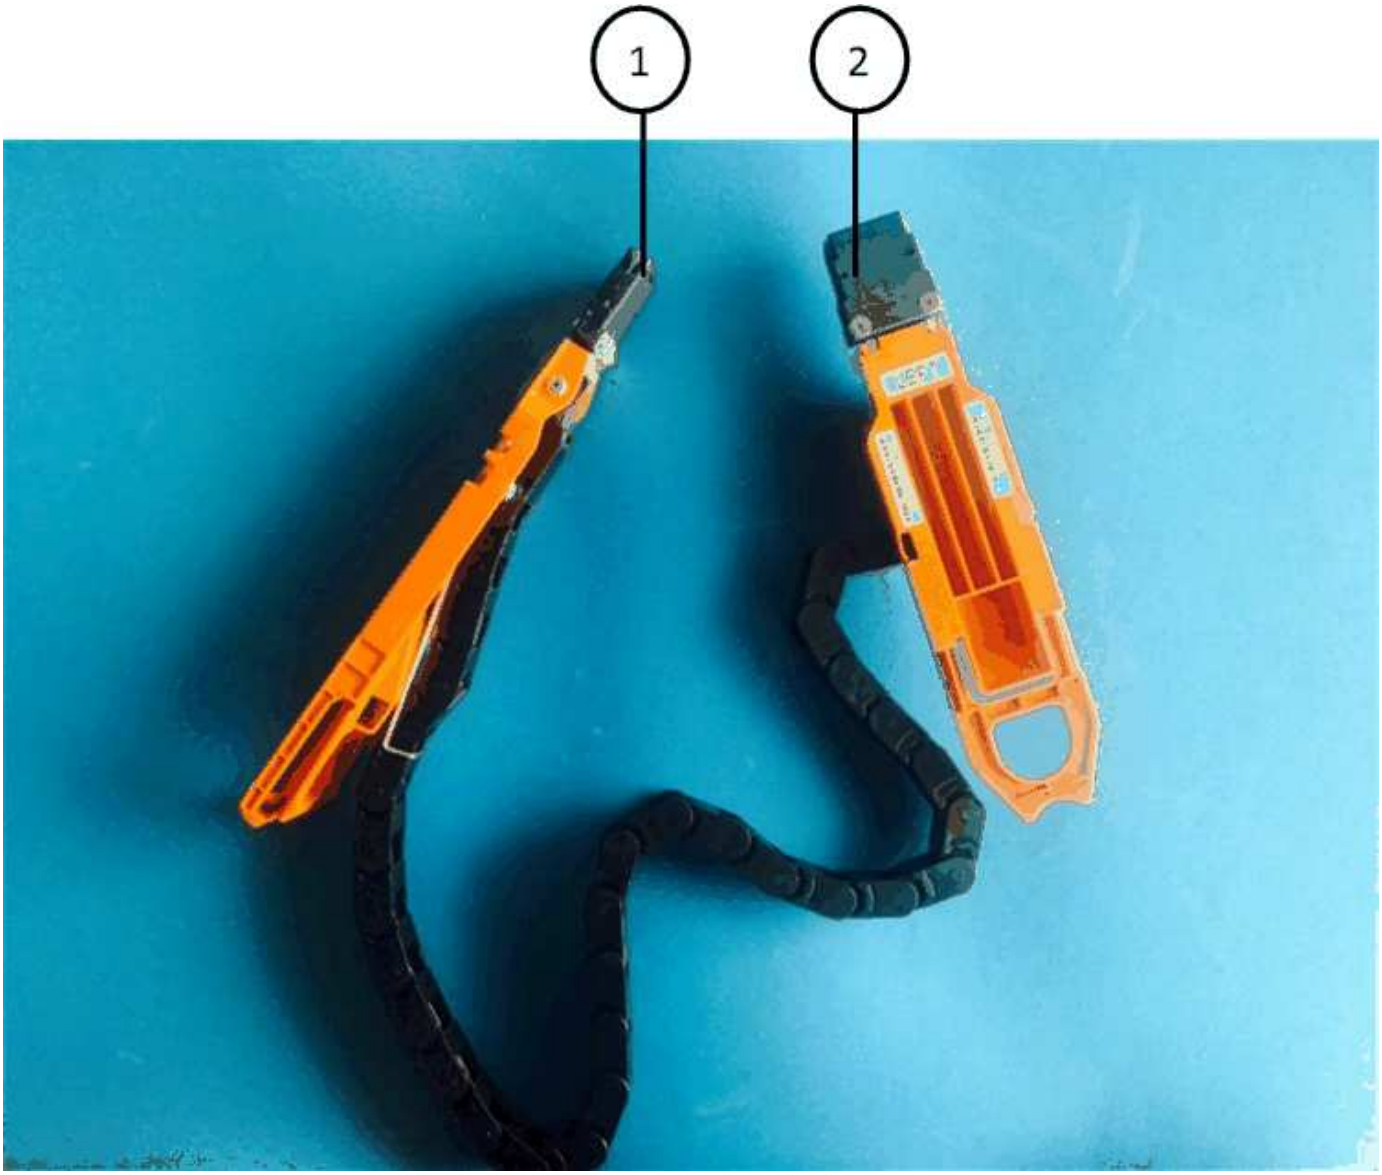

Callout	Cable chain	Connector	Connects to
1	Left	Vertical	Midplane
2	Left	Horizontal	Drive drawer

Callout	Cable chain	Connector	Connects to
1	Right	Horizontal	Drive drawer
2	Right	Vertical	Midplane

Steps

1. Follow these steps to attach the left cable chain:
   a. Locate the horizontal and vertical connectors on the left cable chain and the corresponding horizontal and vertical brackets inside the enclosure.
   b. Align both cable chain connectors with their corresponding brackets.
   c. Slide the cable chain's horizontal connector under the guide rail on the horizontal bracket, and push it in as far as it can go.

      The figure shows the guide rail on the left side for the second drive drawer in the enclosure.

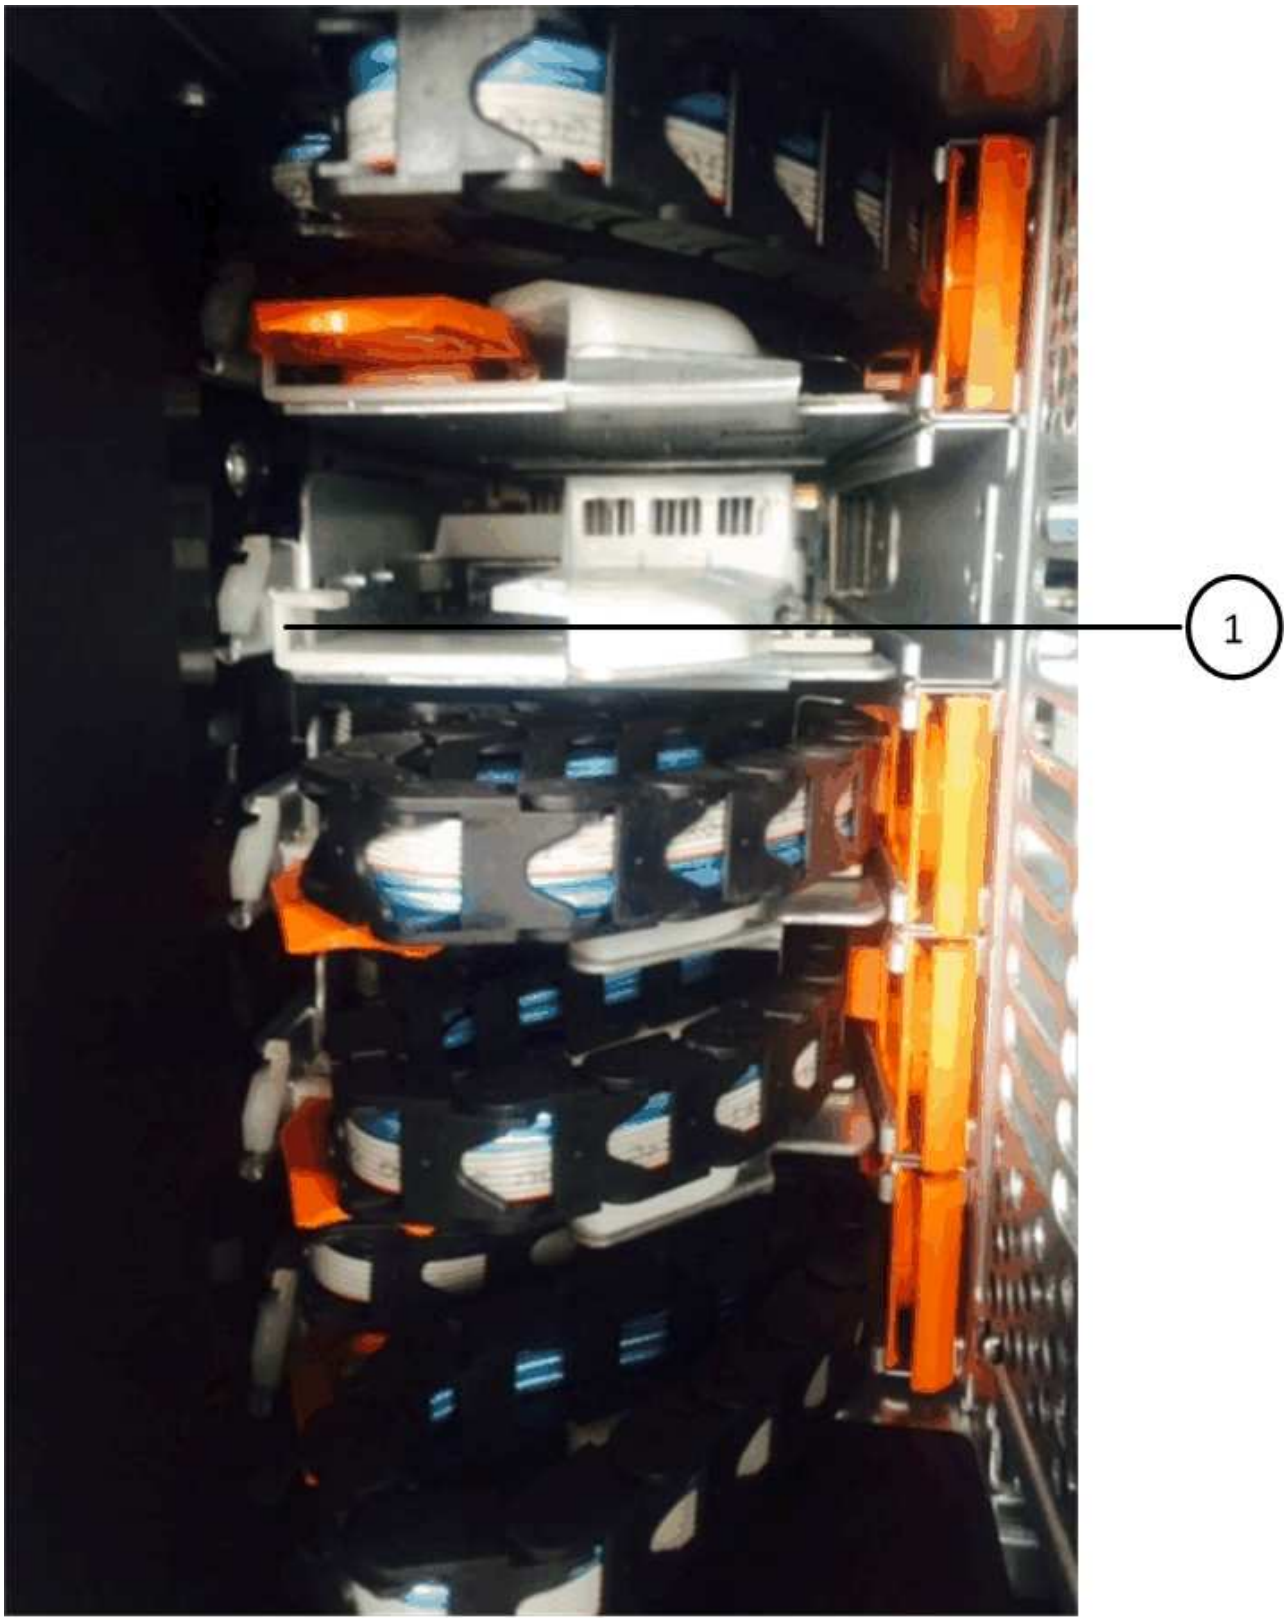

1	Guide rail

Risk of equipment malfunction: Make sure to slide the connector underneath the guide rail on the bracket. If the connector rests on the top of the guide rail, problems might occur when the system runs.

d. Slide the vertical connector on the left cable chain into the vertical bracket.

e. After you have reconnected both ends of the cable chain, carefully pull on the cable chain to verify that both connectors are latched.

Risk of equipment malfunction: If the connectors are not latched, the cable chain might come loose during drawer operation.

2. Reinstall the left fan module.
3. Follow these steps to reattach the right cable chain:
   a. Locate the horizontal and vertical connectors on the cable chain and their corresponding horizontal and vertical brackets inside the enclosure.
   b. Align both cable chain connectors with their corresponding brackets.
   c. Slide the cable chain's horizontal connector under the guide rail on the horizontal bracket and push it in as far as it will go.

Risk of equipment malfunction: Make sure to slide the connector underneath the guide rail on the bracket. If the connector rests on the top of the guide rail, problems might occur when the system runs.

   d. Slide the vertical connector on the right cable chain into the vertical bracket.
   e. After you reconnect both ends of the cable chain, carefully pull on the cable chain to verify that both connectors are latched.

Risk of equipment malfunction: If the connectors are not latched, the cable chain might come loose during drawer operation.

4. Reinstall the right fan module.
5. Reapply power:
   a. Turn on both power switches on the drive shelf.
   b. Confirm that both fans come on and that the amber LED on the back of the fans is off.
6. If you had halted the HA pair, boot ONTAP on both controllers; otherwise go to the next step.
7. If you had moved data off the shelf and deleted the data aggregates, you can now use the spare disks in the shelf for aggregate creation or expansion. To learn more about these procedures, you can refer to the Aggregate creation workflow and Aggregate expansion workflow.

Drive shelf

Overview of shelf maintenance - DS212C, DS224C, or DS460C

You can take the following actions to maintain your SAS shelf:

- Hot-add a drive

- Cold-replace a shelf
- Hot-remove a shelf
- Monitor shelf LEDs

Cold-replace a shelf - DS212C, DS224C, and DS460C

When replacing a drive shelf in a production system that has disks in use, you need to perform a cold shelf replacement. This is a disruptive procedure for shelves with IOM12 or IOM12B modules. It requires you to halt the controllers in your HA pair.

Use the NetApp Knowledge Base article How to replace a shelf chassis using a cold shelf removal procedure.

Hot-add a drive - DS212C, DS224C, or DS460C

You can add new drives to a powered-on shelf non-disruptively, even during I/O operations.

Use the NetApp Knowledge Base article Best practices for adding disks to an existing shelf or cluster.

Hot-remove a shelf - DS212C, DS224C, and DS460C

You can hot-remove a disk shelf with IOM12 or IOM12B modules (nondisruptively remove a disk shelf from a system that is powered on and I/O is in progress) when you need to move or replace a disk shelf. You can hot-remove one or more disk shelves from anywhere within a stack of disk shelves or remove a stack of disk shelves.

About this task

- If you are hot-removing a disk shelf from a stack (but keeping the stack), you recable and verify one path at a time (path A then path B) to bypass the disk shelf you are removing so that you always maintain single-path connectivity from the controllers to the stack.

  If you do not maintain single-path connectivity from the controllers to the stack when recabling the stack to bypass the disk shelf you are removing, you could fail the system with a multidisk panic.

- **Possible shelf damage:** If you are removing a DS460C shelf and you are moving it to a different part of the data center or transporting it to a different location, see the "Move or transport DS460C shelves" section at the end of this procedure.

Before you begin

- **Best practice**: Remove disk drive ownership after you remove the aggregates from the disk drives in the disk shelves you are removing.

  Removing ownership information from a spare disk drive allows the disk drive to be properly integrated into another node (as needed).

  You must disable disk ownership automatic assignment before you remove ownership from disk drives. You reenable this feature at the end of this procedure. To learn more about disk aggregates, see Disks and aggregates overview.

- Your system must be a multipath HA, tri-path HA, multipath, quad-path HA, or quad-path configuration.

  For platforms with internal storage, the external storage must be cabled as multipath HA, tri-path HA, or multipath.

  > For a FAS2600 series single-controller system that has the external storage cabled with multipath connectivity, the system is a mixed-path configuration because the internal storage uses single-path connectivity.

- Your system cannot have any SAS cabling error messages.

  To view any SAS cabling error messages and the corrective actions you should take, download and run the Active IQ Config Advisor.

- HA pair configurations cannot be in a takeover state.
- You must have removed all aggregates from the disk drives (the disk drives must be spares) in the disk shelves you are removing.

  > If you attempt this procedure with aggregates on the disk shelf you are removing, you could fail the system with a multidisk panic.

  You can use the `storage aggregate offline -aggregate` *`aggregate_name`* command and then the `storage aggregate delete -aggregate` *`aggregate_name`* command.

- If you are removing one or more disk shelves from within a stack, you must have factored the distance to bypass the disk shelves you are removing; therefore, if the current cables are not long enough, you need to have longer cables available.
- **Best practice**: For a clustered ONTAP system that is greater than two-nodes, reassign epsilon to an HA pair other than the one that is undergoing planned maintenance.

  Reassigning epsilon minimizes the risk of unforeseen errors impacting all nodes in a clustered ONTAP system. You can use the following steps to determine the node holding epsilon and reassign epsilon if needed:

  1. Set privilege level to advanced: `set -privilege advanced`
  2. Determine which node holds epsilon: `cluster show`

     The node that holds epsilon shows `true` in the `Epsilon` column. (The nodes that do not hold epsilon show `false`.)

  3. If the node in the HA pair that is undergoing maintenance shows `true` (holds epsilon), then remove epsilon from the node: `cluster modify -node` *`node_name`* `-epsilon false`
  4. Assign epsilon to a node in another HA pair: `cluster modify -node` *`node_name`* `-epsilon true`
  5. Return to the admin privilege level: `set -privilege admin`

Steps

1. Verify that your system configuration is `Multi-Path HA`, `tri-path HA`, `Multi-Path`, `Quad-path HA`, or `Quad-path` by running the the `sysconfig` command from the nodeshell of either controller.

It might take up to a minute for the system to complete discovery.

The configuration is listed in the `System Storage Configuration` field.

> For a FAS2600 series single-controller system that has the external storage cabled with multipath connectivity, the output is displayed as `mixed-path` because the internal storage uses single-path connectivity.

2. Verify that the disk drives in the disk shelves you are removing have no aggregates (are spares) and ownership is removed:

   a. Enter the following command from the clustershell of either controller: `storage disk show -shelf` *`shelf_number`*

   b. Check the output to verify that there are no aggregates on the disk drives in the disk shelves you are removing.

      Disk drives with no aggregates have a dash in the `Container Name` column.

   c. Check the output to verify that ownership is removed from the disk drives on the disk shelves you are removing.

      Disk drives with no ownership have a dash in the `Owner` column.

      > If you have failed disk drives in the shelf you are removing, they have broken in the `Container Type` column. (Failed disk drive do not have ownership.)

      The following output shows disk drives on the disk shelf being removed (disk shelf 3) are in a correct state for removing the disk shelf. The aggregates are removed on all of the disk drives; therefore, a dash appears in the `Container Name` column for each disk drive. Ownership is also removed on all of the disk drives; therefore, a dash appears in the `Owner` column for each disk drive.

```
cluster::> storage disk show -shelf 3

           Usable           Disk   Container   Container
Disk         Size Shelf Bay Type   Type        Name       Owner
-------- -------- ----- --- ------ ----------- ---------- ---------
...
1.3.4           -     3   4 SAS    spare                 -         -
1.3.5           -     3   5 SAS    spare                 -         -
1.3.6           -     3   6 SAS    broken                -         -
1.3.7           -     3   7 SAS    spare                 -         -
...
```

3. Physically locate the disk shelves you are removing.

   If needed, you can turn on the disk shelf's location (blue) LEDs to aid in physically locating the affected disk shelf: `storage shelf location-led modify -shelf-name` *`shelf_name`* `-led-status on`

A disk shelf has three location LEDs: one on the operator display panel and one on each IOM12 module. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the off option.

4. If you are removing an entire stack of disk shelves, complete the following substeps; otherwise, go to the next step:
   a. Remove all SAS cables on path A (IOM A) and path B (IOM B).

      This includes controller-to-shelf cables and shelf-to-shelf cables for all disk shelves in the stack you are removing.

   b. Go to step 9.
5. If you are removing one or more disk shelves from a stack (but keeping the stack), recable the path A (IOM A) stack connections to bypass the disk shelves you are removing by completing the applicable set of substeps:

   If you are removing more than one disk shelf in the stack, complete the applicable set of substeps one disk shelf at a time.

Wait at least 10 seconds before connecting the port. The SAS cable connectors are keyed; when oriented correctly into a SAS port, the connector clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector).

If you are removing…	Then…
A disk shelf from either end (logical first or last disk shelf) of a stack	a. Remove any shelf-to-shelf cabling from IOM A ports on the disk shelf you are removing and put them aside. b. Unplug any controller-to-stack cabling connected to IOM A ports on the disk shelf you are removing and plug them into the same IOM A ports on the next disk shelf in the stack. The “next” disk shelf can be above or below the disk shelf you are removing depending on which end of the stack you are removing the disk shelf from.

If you are removing…	Then…
A disk shelf from the middle of the stack A disk shelf in the middle of a stack is only connected to other disk shelves—not to any controllers.	a. Remove any shelf-to-shelf cabling from IOM A ports 1 and 2 or from ports 3 and 4 on the disk shelf you are removing and IOM A of the next disk shelf, and then put them aside. b. Unplug the remaining shelf-to-shelf cabling connected to IOM A ports on the disk shelf you are removing and plug them into the same IOM A ports on the next disk shelf in the stack. The “next” disk shelf can be above or below the disk shelf you are removing depending on which IOM A ports (1 and 2 or 3 and 4) you removed the cabling from.

You can refer to the following cabling examples when removing a disk shelf from an end of a stack or the middle of a stack. Note the following about the cabling examples:

- The IOM12/IOM12B modules are arranged side-by-side as in a DS224C or DS212C disk shelf; if you have a DS460C, the IOM12/IOM12B modules are arranged one above the other.
- The stack in each example is cabled with standard shelf-to-shelf cabling, which is used in stacks cabled with multipath HA, tri-path HA, or multipath connectivity.

  You can infer the recabling if your stack is cabled with quad-path HA or quad-path connectivity, which uses double-wide shelf-to-shelf cabling.

- The cabling examples demonstrate recabling one of the paths: path A (IOM A).

  You repeat the recabling for path B (IOM B).

- The cabling example for removing a disk shelf from the end of a stack demonstrates removing the logical last disk shelf in a stack that is cabled with multipath HA or tri-path HA connectivity.

  You can infer the recabling if you are removing the logical first disk shelf in a stack or if your stack has multipath connectivity.

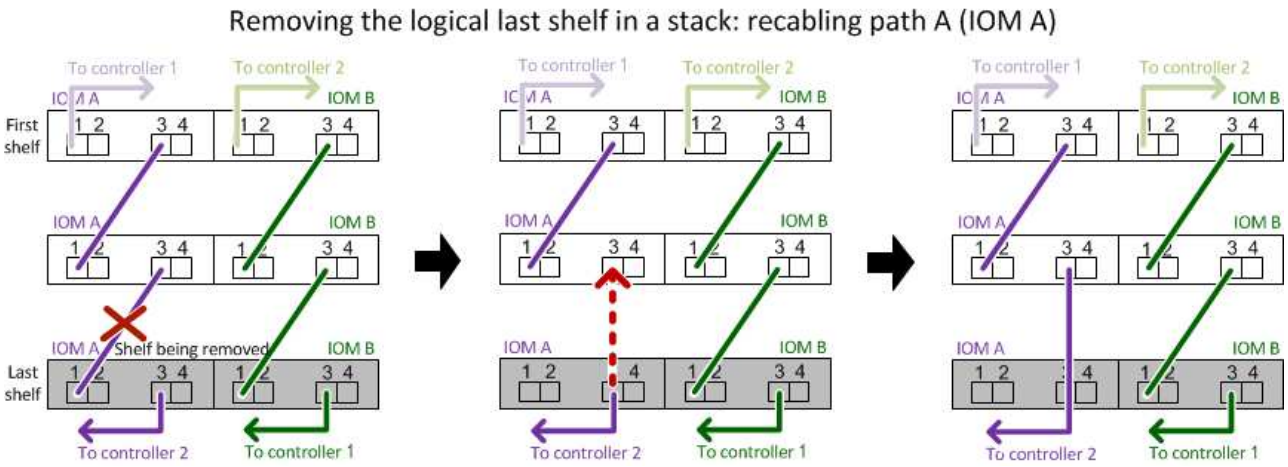

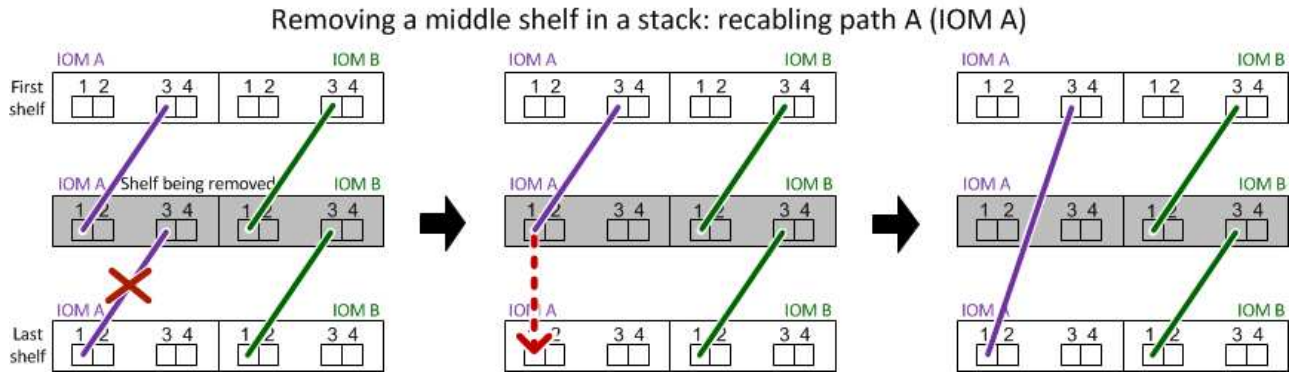

6. Verify that you bypassed the disk shelves you are removing and reestablished the path A (IOM A) stack connections correctly: `storage disk show -port`

   For HA pair configurations, you run this command from the clustershell of either controller. It might take up to a minute for the system to complete discovery.

   The first two lines of output show disk drives with connectivity through both path A and path B. The last two lines of output show disk drives with connectivity through a single-path, path B.

```
cluster::> storage show disk -port

PRIMARY   PORT SECONDARY      PORT TYPE SHELF BAY
--------- ---- ---------      ---- ---- ----- ---
1.20.0    A    node1:6a.20.0  B    SAS  20    0
1.20.1    A    node1:6a.20.1  B    SAS  20    1
1.21.0    B    -              -    SAS  21    0
1.21.1    B    -              -    SAS  21    1
...
```

7. The next step depends on the `storage disk show -port` command output:

If the output shows…	Then…
All disk drives in the stack are connected through path A and path B except for the ones in the disk shelves you disconnected, which are only connected through path B	Go to the next step. You successfully bypassed the disk shelves you are removing and reestablished path A on the remaining disk drives in the stack.
Anything other than the above	Repeat Step 5 and Step 6. You must correct the cabling.

8. Complete the following substeps for the disk shelves (in the stack) you are removing:
   a. Repeat Step 5 through Step 7 for path B.

When you repeat Step 7 and if you have recabled the stack correctly, you should only see all remaining disk drives connected through path A and path B.

b. Repeat Step 1 to confirm that your system configuration is the same as before you removed one or more disk shelves from a stack.

c. Go to the next step.

9. If when you removed ownership from the disk drives (as part of the preparation for this procedure), you disabled disk ownership automatic assignment, reenable it by entering the following command; otherwise, go to the next step: `storage disk option modify -autoassign on`

   For HA pair configurations, you run the command from the clustershell of both controllers.

10. Power off the disk shelves you disconnected and unplug the power cords from the disk shelves.

11. Remove the disk shelves from the rack or cabinet.

    To make a disk shelf lighter and easier to maneuver, remove the power supplies and I/O modules (IOMs).

    For DS460C disk shelves, a fully loaded shelf can weigh approximately 247 lbs (112 kg); therefore, exercise the following caution when removing a shelf from a rack or cabinet.

    It is recommended that you use a mechanized lift or four people using the lift handles to safely move a DS460C shelf.

    Your DS460C shipment was packaged with four detachable lift handles (two for each side). To use the lift handles, you install them by inserting the tabs of the handles into the slots in the side of the shelf and pushing up until they click into place. Then, as you slide the disk shelf onto the rails, you detach one set of handles at a time using the thumb latch. The following illustration shows how to attach a lift handle.

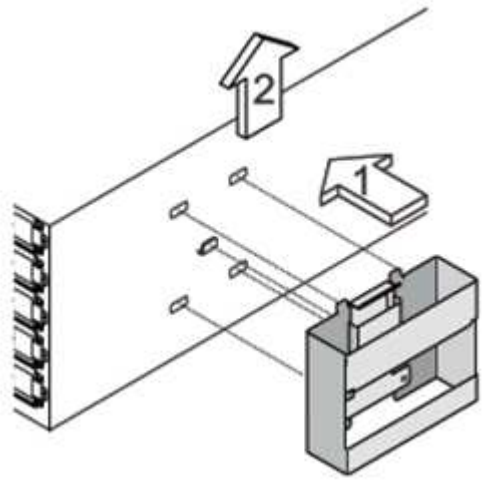

    If you are moving the DS460C shelf to a different part of the data center or transporting it to a different location, see the "Move or transport DS460C shelves" section.

Move or transport DS460C shelves

If you move a DS460C shelf to a different part of the data center or transport the shelf to a different location, you need to remove the drives from the drive drawers to avoid possible damage to the drive drawers and drives.

- If when you installed DS460C shelves as part of your new system installation or shelf hot-add, you saved the drive packaging materials, use these to repackage the drives before moving them.

  If you did not save the packaging materials, you should place drives on cushioned surfaces or use alternate cushioned packaging. Never stack drives on top of each other.

- Before handling drives, wear an ESD wrist strap grounded to an unpainted surface on your storage enclosure chassis.

  If a wrist strap is unavailable, touch an unpainted surface on your storage enclosure chassis before handling a drive.

- You should take steps to handle drives carefully:
  - Always use two hands when removing, installing, or carrying a drive to support its weight.

    Do not place hands on the drive boards exposed on the underside of the drive carrier.

  - Be careful not to bump drives against other surfaces.
  - Drives should be kept away from magnetic devices.

    Magnetic fields can destroy all data on a drive and cause irreparable damage to the drive circuitry.

Monitor shelf LEDs - DS212C, DS224C, or DS460C

You can monitor the health of your disk shelf by understanding the location and status conditions of the LEDs on your disk shelf components.

Operator display panel LEDs

The LEDs on the disk shelf front operator display panel indicate whether your disk shelf is functioning normally or there are problems with the hardware.

The following table describes the three LEDs on the operator display panel used in DS460C, DS224C, and DS212C disk shelves:

LED icon	LED name	State	Description
	Power	Solid green	One or more power supplies are supplying power to the disk shelf.

LED icon	LED name	State	Description
!	Attention	Solid amber	An error occurred with the function of one of more FRUs: the disk shelf, disk drives, IOM12/IOM12B modules, or power supplies. Check event messages to determine corrective action to take.
		Blinking amber	The shelf ID is in a pending state. Power cycle the disk shelf for the shelf ID to take affect.
	Location	Solid blue	The system administrator activated this LED function to aid in physically locating the disk shelf requiring service. The location LED on the operator display panel and both IOM12/IOM12B modules illuminate when this LED function is activated. Location LEDs automatically turn off after 30 minutes.

Depending on your disk shelf model, the operator display panel looks different; however, the three LEDs are arranged in the same way.

The following illustration is of a DS224C disk shelf operator display panel with the end cap on:

IOM12/IOM12B module LEDs

The LEDs on the IOM12/IOM12B module indicate whether the module is functioning normally, whether it is ready for I/O traffic, and whether there are any problems with the hardware.

The following table describes IOM12/IOM12B module LEDs associated with the function of the module and the function of each SAS port on the module.

The IOM12/IOM12B module is used in DS460C, DS224C, and DS212C disk shelves.

LED icon	LED name	State	Description
!	Attention	Solid amber	IOM12/IOM12B module function: An error occurred with the function of the IOM12/IOM12B module. SAS port function: Less than all four SAS lanes established a link (with either an adapter or another disk shelf). Check event messages to determine corrective action to take.
LNK	Port link	Solid green	One or more of the four SAS lanes established a link (with either an adapter or another disk shelf).

LED icon	LED name	State	Description
	Location	Solid blue	The system administrator activated this LED function to aid in physically locating the disk shelf with the failed IOM12/IOM12B module. The location LED on the operator display panel and both IOM12/IOM12B modules illuminate when this LED function is activated. Location LEDs automatically turn off after 30 minutes.

The following illustration is for an IOM12 module:

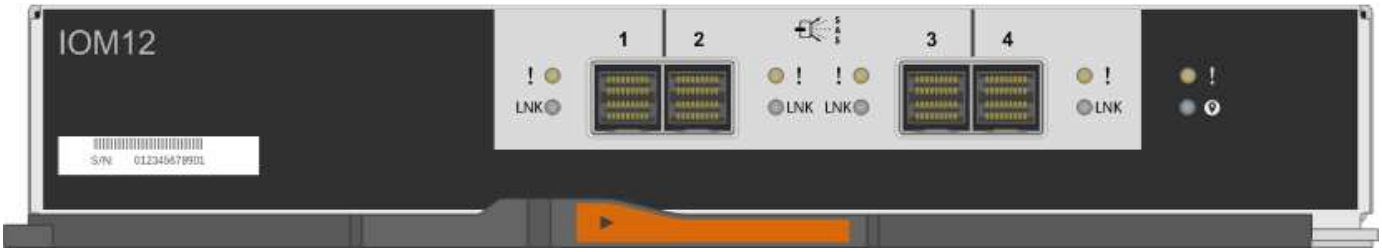

The IOM12B modules are distinguished by a blue stripe and an "IOM12B" label:

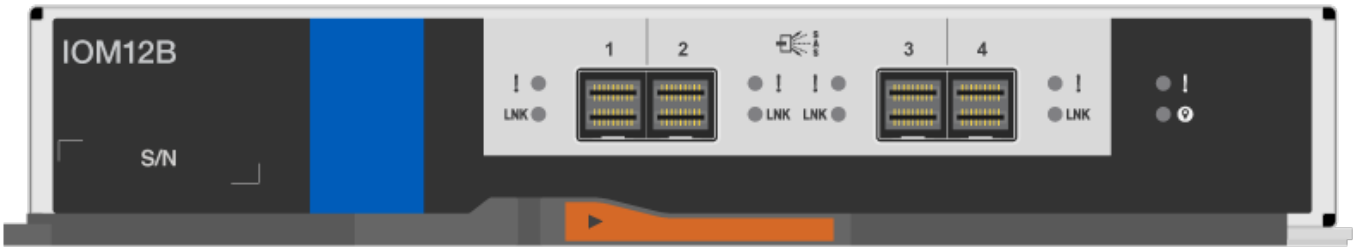

Power supply LEDs

The LEDs on the power supply indicate whether the power supply is functioning normally or there are hardware problems.

The following table describes the two LEDs on power supplies used in DS460C, DS224C, and DS212C disk shelves:

LED icon	LED name	State	Description
	Power	Solid green	The power supply is functioning correctly.
		Off	The power supply failed, the AC switch is turned off, the AC power cord is not properly installed, or electricity is not being properly supplied to the power supply. Check event messages to determine corrective action to take.
!	Attention	Solid amber	An error occurred with the function of the power supply. Check event messages to determine corrective action to take.

Depending on your disk shelf model, power supplies can be different, dictating the location of the two LEDs.

The following illustration is for a power supply used in a DS460C disk shelf.

The two LED icons act as the labels and LEDs, meaning the icons themselves illuminate—there are no adjacent LEDs.

The following illustration is for a power supply used in a DS224C or DS212C disk shelf:

Fan LEDs on DS460C disk shelves

The LEDs on the DS460C fans indicate whether the fan is functioning normally or there are hardware problems.

The following table describes the LEDs on fans used in DS460C disk shelves:

Item	LED name	State	Description
1	Attention	Solid amber	An error occurred with the function of the fan. Check event messages to determine corrective action to take.

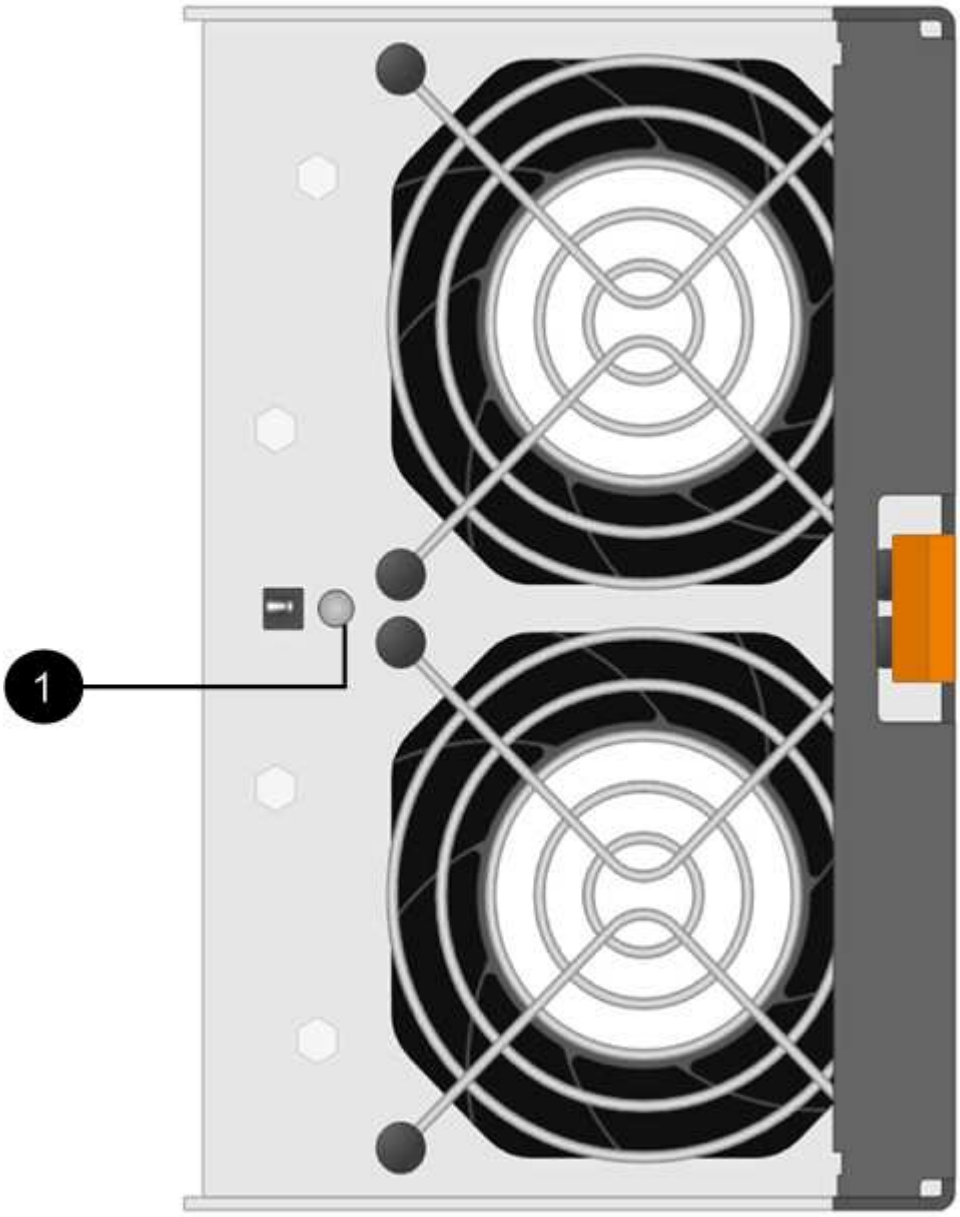

Disk drive LEDs

The LEDs on a disk drive indicates whether it is functioning normally or there are problems with the hardware.

Disk drive LEDs for DS224C and DS212C disk shelves

The following table describes the two LEDs on the disk drives used in DS224C and DS212C disk shelves:

Callout	LED name	State	Description
1	Activity	Solid green	The disk drive has power.
		Blinking green	The disk drive has power and I/O operations are in progress.
2	Attention	Solid amber	An error occurred with the function of the disk drive. Check event messages to determine corrective action to take.

Depending on your disk shelf model, disk drives are arranged vertically or horizontally in the disk shelf, dictating the location of the two LEDs.

The following illustration is for a disk drive used in a DS224C disk shelf.

DS224C disk shelves use 2.5-inch disk drives arranged vertically in the disk shelf.

The following illustration is for a disk drive used in a DS212C disk shelf.

DS212C disk shelves use 3.5-inch disk drives or 2.5-inch disk drives in carriers arranged horizontally in the disk shelf.

Disk drive LEDs for DS460C disk shelves

The following illustration and table describes the drive activity LEDs on the drive drawer and their operational states:

Location	LED	Status indicator	Description
1	Attention: Drawer attention for each drawer	Solid amber	A component within the drive drawer requires operator attention.
		Off	No drive or other component in the drawer requires attention and no drive in the drawer has an active locate operation.
		Blinking amber	A locate drive operation is active for any drive within the drawer.
2-13	Activity: Drive activity for drives 0 through 11 in the drive drawer	Green	The power is turned on and the drive is operating normally.
		Blinking green	The drive has power, and I/O operations are in progress.
		Off	The power is turned off.

When the drive drawer is open, an attention LED can be seen in front of each drive.

1	Attention LED light on

Replace a fan module - DS460C

Each DS460C drive shelf includes two fan modules. If a fan module fails, you must replace it as soon as possible to ensure that the shelf has adequate cooling. When you remove the failed fan module, you do not have to turn off power to your disk shelf.

This procedure applies to shelves with IOM12 or IOM12B modules.

Before you begin

You must ensure that you remove and replace the fan module within 30 minutes to prevent the system from overheating.

Steps

1. Put on antistatic protection.
2. Unpack the new fan module, and place it on a level surface near the shelf.

   Save all packing material for use when returning the failed fan.

3. From the back of the disk shelf, look at the Attention LEDs to locate the fan module you need to remove.

   You must replace the fan module that has its Attention LED on.

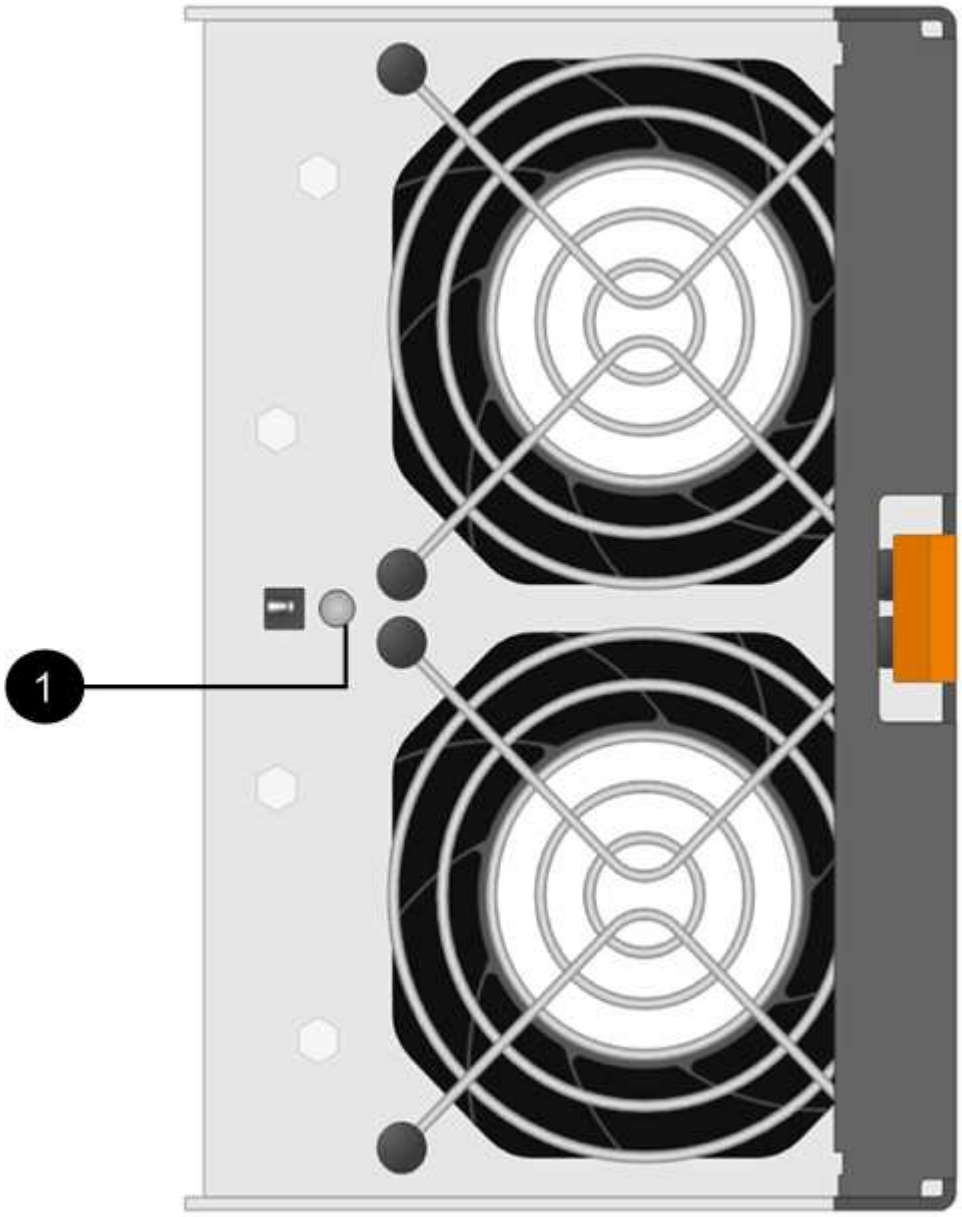

Item	LED name	State	Description
1	Attention	Solid amber	The fan has a fault

4. Press the orange tab to release the fan module handle.

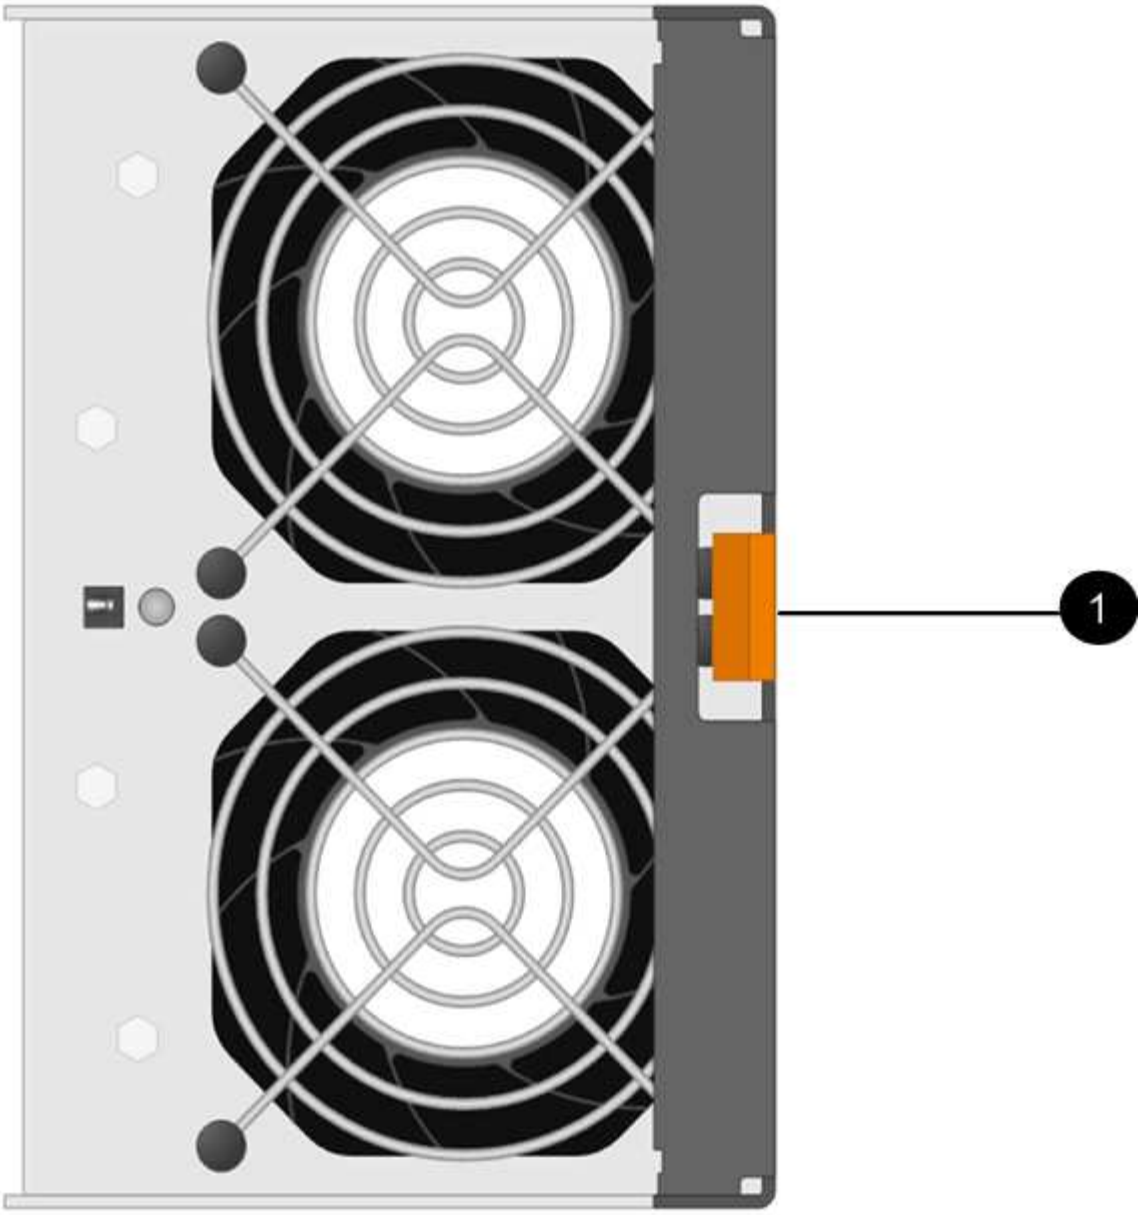

1	Fan module release tab

5. Use the fan module handle to pull the fan module out of the shelf.

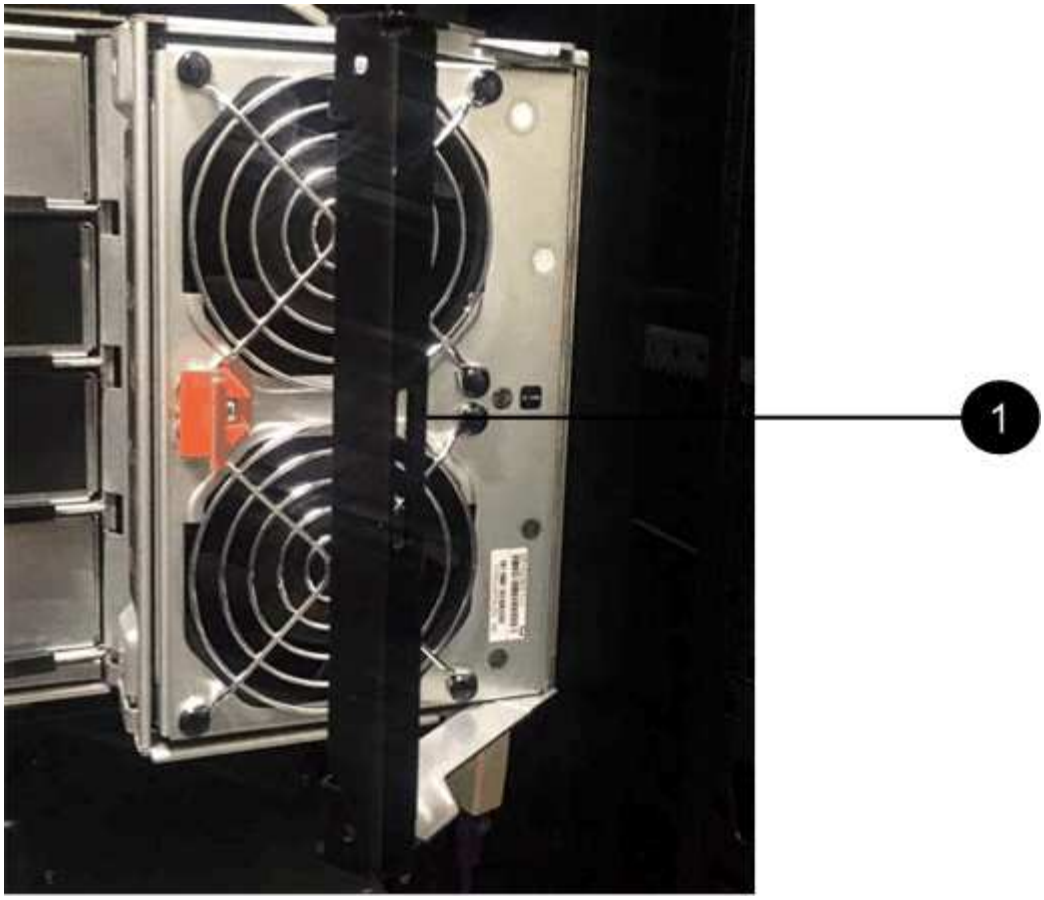

1	Handle to pull the fan module out

6. Slide the replacement fan module all the way into the shelf, moving the fan module handle to the side until it latches with the orange tab.
7. Check the amber Attention LED on the new fan module.

> After you replace the fan module, the Attention LED stays on (solid amber) while the firmware checks that the fan module was installed correctly. The LED goes off after this process is complete.

8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

   Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number.

Hot-swap or replace an IOM module - DS212C, DS224C, or DS460C

Your system configuration determines whether you can perform a nondisruptive shelf IOM hot-swap or a disruptive shelf IOM replacement when a IOM12 or IOM12B shelf IOM fails.

About this task

- This procedure applies to shelves having IOM12 or IOM12B modules.

This procedure is for like-for-like shelf IOM hot-swaps or replacements. This means you can only replace an IOM12 module with another IOM12 module or an IOM12B module with another IOM12B module.

- The IOM12 or IOM12B modules can be distinguished by their appearance:

  The IOM12 modules are distinguished by an "IOM12" label:

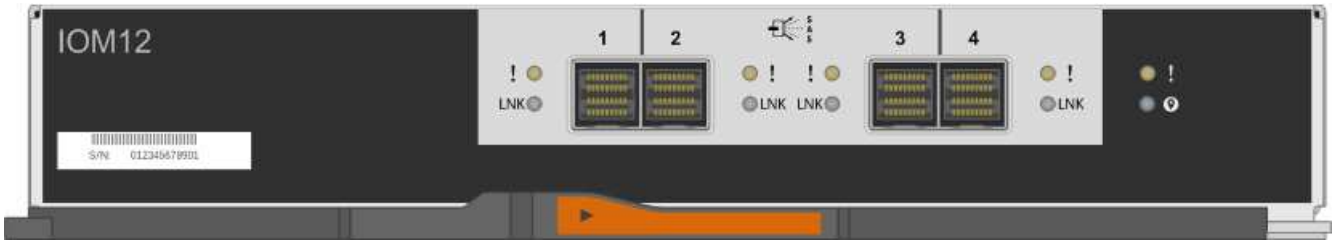

  The IOM12B modules are distinguished by a blue stripe and an "IOM12B" label:

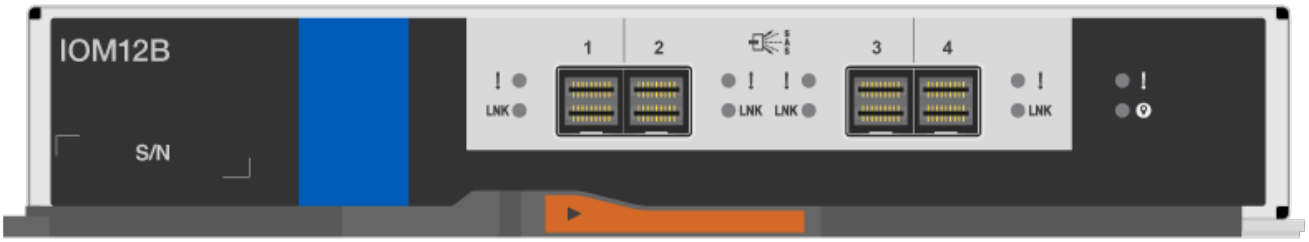

- For multipathed (multipath HA or multipath), tri-path HA, and quad-pathed (quad-path HA or quad-path) configurations, you can hot-swap a shelf IOM (nondisruptively replace a shelf IOM in a system that is powered on and serving data—I/O is in progress).
- For FAS2600 series and FAS2700 series single-path HA configurations, you must perform a takeover and giveback operation to replace a shelf IOM in a system that is powered on and serving data—I/O is in progress.
- For FAS2600 series single-path configurations, you must halt your system to replace a shelf IOM.

  If you attempt to hot-swap a shelf IOM on a disk shelf with a single-path connection, you will lose all access to the disk drives in the disk shelf as well as any disk shelves beneath. You could also bring down your entire system.

- Disk shelf (IOM) firmware is automatically updated (nondisruptively) on a new shelf IOM with a non-current firmware version.

  Shelf IOM firmware checks occur every ten minutes. An IOM firmware update can take up to 30 minutes.

- If needed, you can turn on the disk shelf's location (blue) LEDs to aid in physically locating the affected disk shelf: `storage shelf location-led modify -shelf-name shelf_name -led-status on`

  A disk shelf has three location LEDs: one on the operator display panel and one on each shelf IOM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the off option.

- If needed, you can refer to the Monitoring disk shelf LEDs guide for information about the meaning and location of disk shelf LEDs on the operator display panel and FRU components.

Before you begin

- All other components in the system—including the other IOM12/IOM12B module—must be functioning properly.
- **Best practice**: Ensure your system has the current versions of disk shelf (IOM) firmware and disk drive firmware before adding new disk shelves, shelf FRU components, or SAS cables. You can visit the NetApp Support Site to download disk shelf firmware and download disk drive firmware.

Steps

1. Properly ground yourself.
2. Unpack the new shelf IOM and set it on a level surface near the disk shelf.

   Save all packaging materials for use when returning the failed shelf IOM.
3. Physically identify the failed shelf IOM from the system console warning message and the illuminated attention (amber) LED on the failed shelf IOM.
4. Perform one of the following actions based on the type of configuration you have:

If you have a…	Then…
Multipath HA, tri-path HA, multipath, quad-path HA, or quad-path configuration	Go to the next step.
FAS2600 series and FAS2700 series single-path HA configuration	a. Determine the target node (the node the failed shelf IOM belongs to). IOM A belongs to Controller 1. IOM B belongs to Controller 2. b. Take over the target node: `storage failover takeover -bynode` *`partner HA node`*
FAS2600 series single-path configuration	a. Shut down the system from the system console: `halt` b. Verify that your system halted by checking the storage system console.

5. Disconnect the cabling from the shelf IOM that you are removing.

   Make note of the shelf IOM ports each cable is connected to.
6. Press the orange latch on the shelf IOM cam handle until it releases, and then open the cam handle fully to release the shelf IOM from the mid plane.

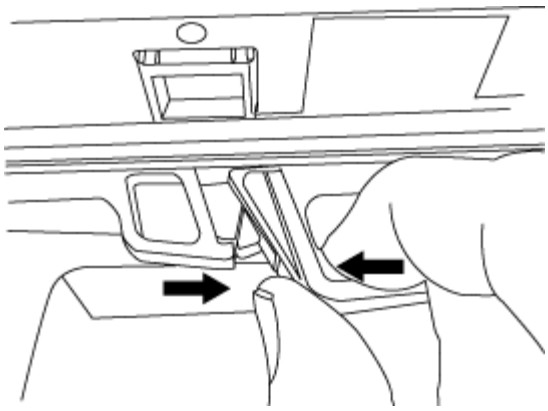

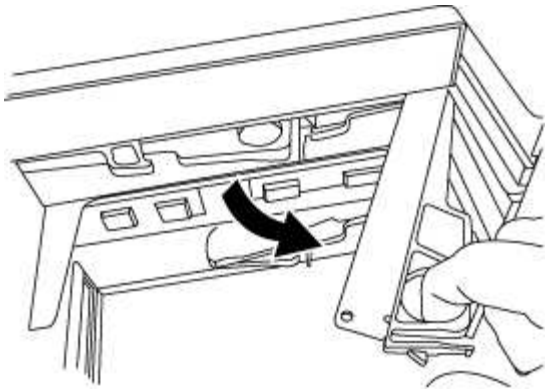

7. Use the cam handle to slide the shelf IOM out of the disk shelf.

   When handling a shelf IOM, always use two hands to support its weight.

8. Wait at least 70 seconds after removing the shelf IOM before you install the new shelf IOM.

   Waiting at least 70 seconds enables the driver to register the shelf ID correctly.

9. Using two hands, with the cam handle of the new shelf IOM in the open position, support and align the edges of the new shelf IOM with the opening in the disk shelf, and then firmly push the new shelf IOM until it meets the mid plane.

   Do not use excessive force when sliding the shelf IOM into the disk shelf; you might damage the connectors.

10. Close the cam handle so that the latch clicks into the locked position and the shelf IOM is fully seated.
11. Reconnect the cabling.

    The SAS cable connectors are keyed; when oriented correctly into an IOM port, the connector clicks into place and the IOM port LNK LED illuminates green. You insert a SAS cable connector into an IOM port with the pull tab oriented down (on the underside of the connector).

12. Perform one of the following actions based on the type of configuration you have:

If you have a…	Then…
Multipath HA, tri-path HA, multipath, quad-path HA, or quad-path configuration	Go to the next step.

If you have a…	Then…
FAS2600 series and FAS2700 series single-path HA configuration	Give back the target node: `storage failover giveback -fromnode partner_HA_node`
FAS2600 series single-path configuration	Reboot your system.

13. Verify that the shelf IOM port links have been established.

    For each module port that you cabled, the LNK (green) LED illuminates when one or more of the four SAS lanes have established a link (with either an adapter or another disk shelf).

14. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Hot-swap a power supply - DS212C, DS224C, or DS460C

You can hot-swap a failed power supply in a DS460C, DS224C, or DS212C disk shelf.

This procedure applies to shelves with IOM12 or IOM12B modules.

Before you begin

- All other components in the system—including the other power supply—must be functioning properly.
- If you are replacing more than one power supply, you must do so one at a time so that the disk shelf maintains power.
- You must replace a power supply within two minutes of removal to minimize disruption to the disk shelf's airflow.
- Always use two hands when removing, installing, or carrying a power supply to support its weight.
- **Best practice**: Ensure your system has the current versions of disk shelf (IOM) firmware and disk drive firmware before adding new disk shelves, shelf FRU components, or SAS cables. You can visit the NetApp Support Site to download disk shelf firmware and download disk drive firmware.
- If needed, you can turn on the disk shelf's location (blue) LEDs to aid in physically locating the affected disk shelf: `storage shelf location-led modify -shelf-name shelf_name -led-status on`

  A disk shelf has three location LEDs: one on the operator display panel and one on each shelf IOM. Location LEDs remain illuminated for 30 minutes. You can turn them off by entering the same command, but using the off option.

- If needed, you can refer to the Monitoring disk shelf LEDs guide for information about the meaning and location of disk shelf LEDs on the operator display panel and FRU components.

Steps

1. Properly ground yourself.
2. Unpack the new power supply and set it on a level surface near the shelf.

   Save all packing materials for use when returning the failed power supply.

3. Physically identify the failed power supply from the system console warning message and the illuminated attention (amber) LED on the power supply.
4. Turn off the failed power supply and disconnect the power cable:
   a. Turn off the power switch on the power supply.
   b. Open the power cord retainer and unplug the power cord from the power supply.
   c. Unplug the power cord from the power source.
5. Press the orange latch on the power supply cam handle until it releases, and then open the cam handle to fully release the power supply from the mid plane.

   The following illustration is for a power supply used in a DS224C or DS212C disk shelf; however, the latch operates the same way for power supplies used in DS460C disk shelves.

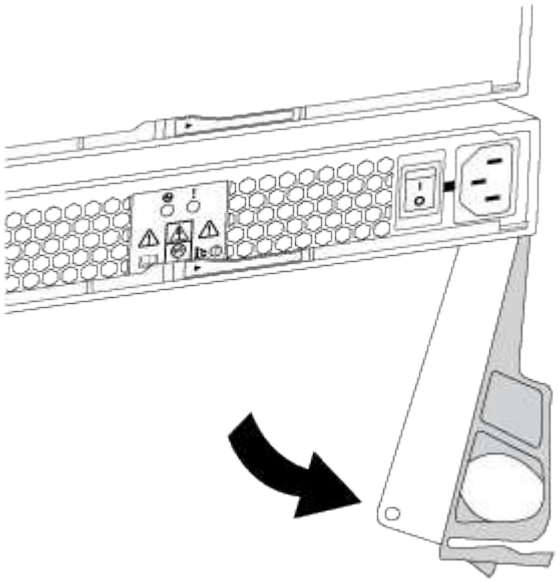

6. Use the cam handle to slide the power supply out of the disk shelf.

   If you have a DS224C or DS212C disk shelf, as you remove the power supply, a flap swings into place to block the empty bay. This helps maintain air flow and cooling.

   When handling a power supply, always use two hands to support its weight.

7. Make sure that the on/off switch of the new power supply is in the Off position.
8. With the cam handle of the new power supply in the open position, use two hands to support and align the edges of the new power supply with the opening in the disk shelf. Then, firmly push the new power supply until it meets the mid plane.

   Do not use excessive force when sliding the power supply into the disk shelf; you might damage the connectors.

9. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.

10. Reconnect the power supply cable and turn on the new power supply:
    a. Reconnect the power cord to the power source.
    b. Reconnect the power cord to the power supply and secure the power cord with the power cord retainer.
    c. Turn on the power switch.

       The power supply's power (green) LED and attention (amber) LED illuminate. Within 40 seconds, the attention (amber) LED turns off.

11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Cabinet and rail kits

SuperRail kit installation instructions

The SuperRail can be either installed on a standard square-hole four-post rack or a standard round-hole four-post rack by using the round-to-square hole adapter brackets.

> The minimum depth between front and rear connection points for installing SuperRail is 24 inches and while the maximum depth is 32 inches.

Installing SuperRail to square-hole four-post rack

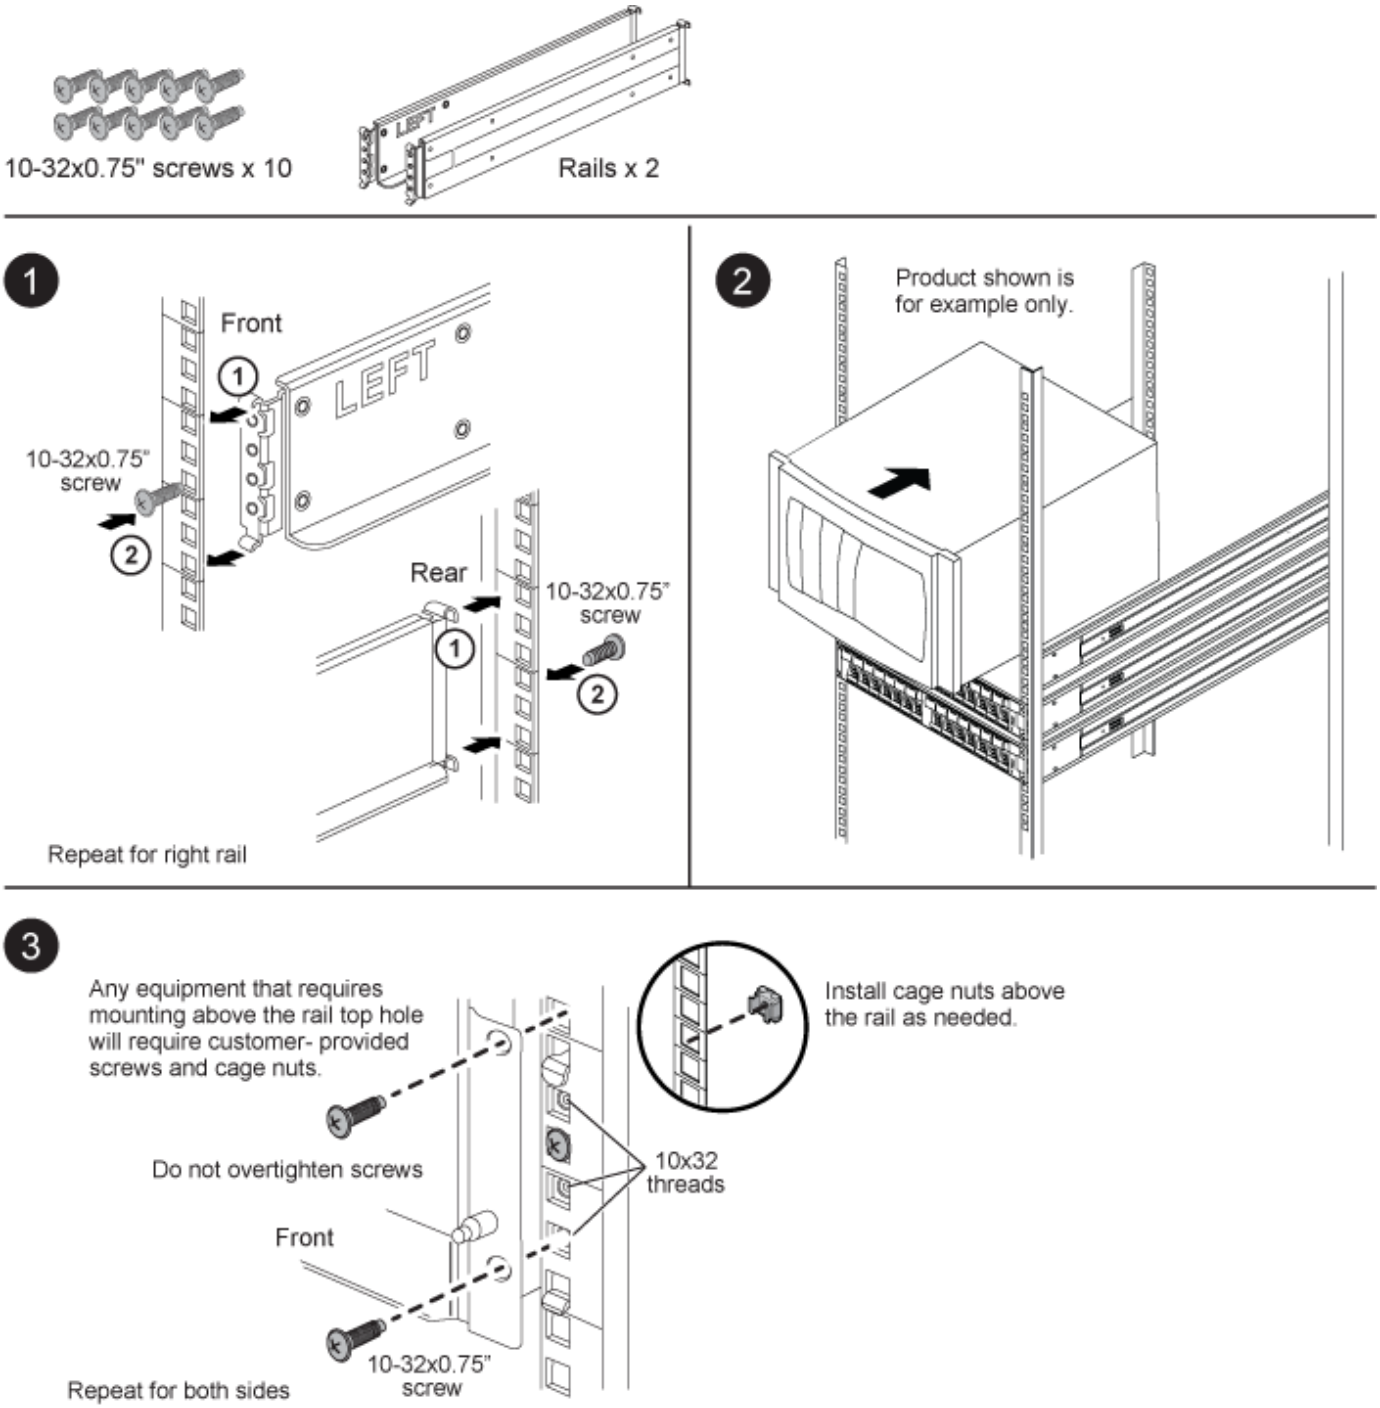

Installing SuperRail to round-hole four-post rack

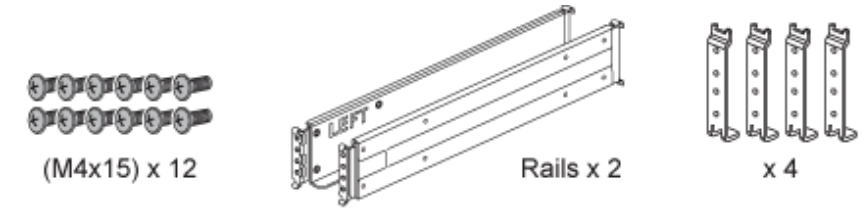

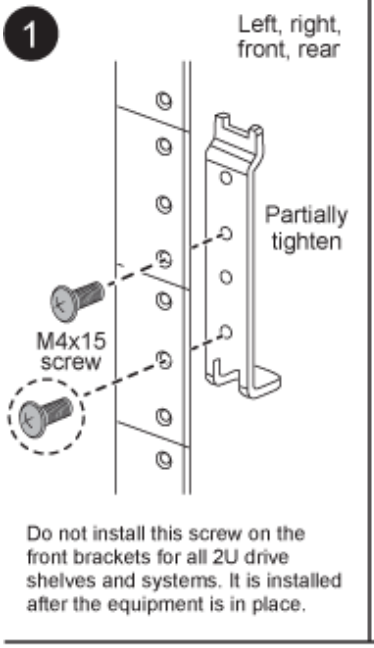

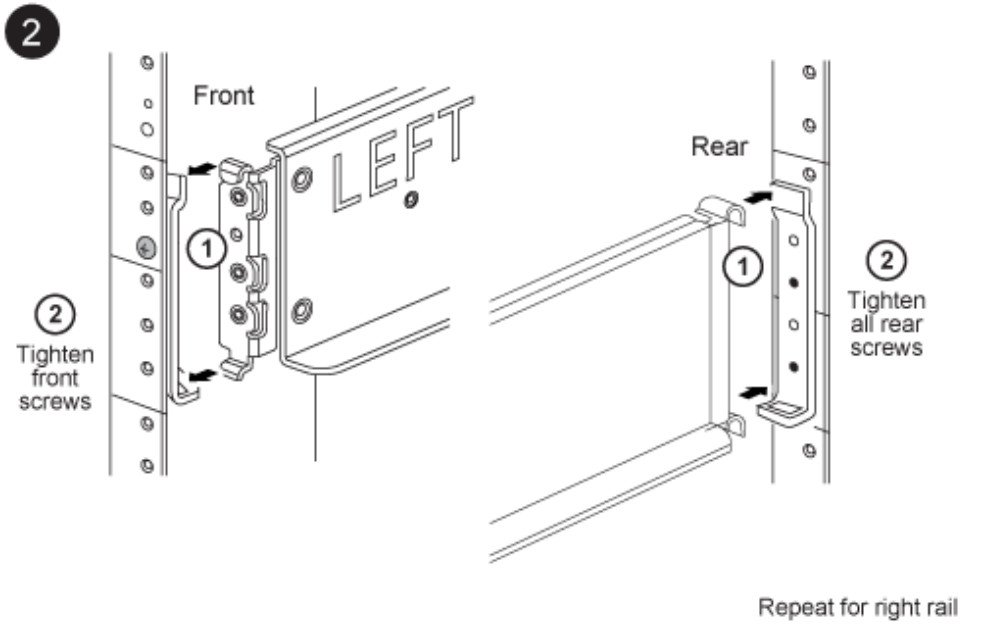

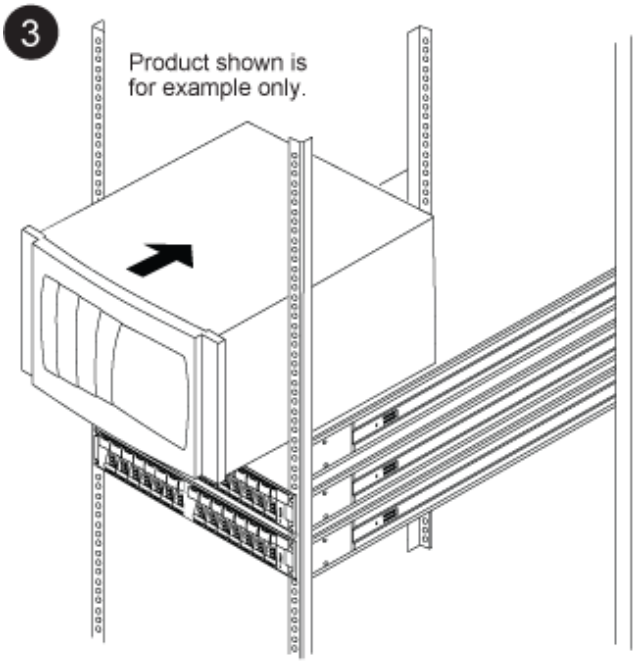

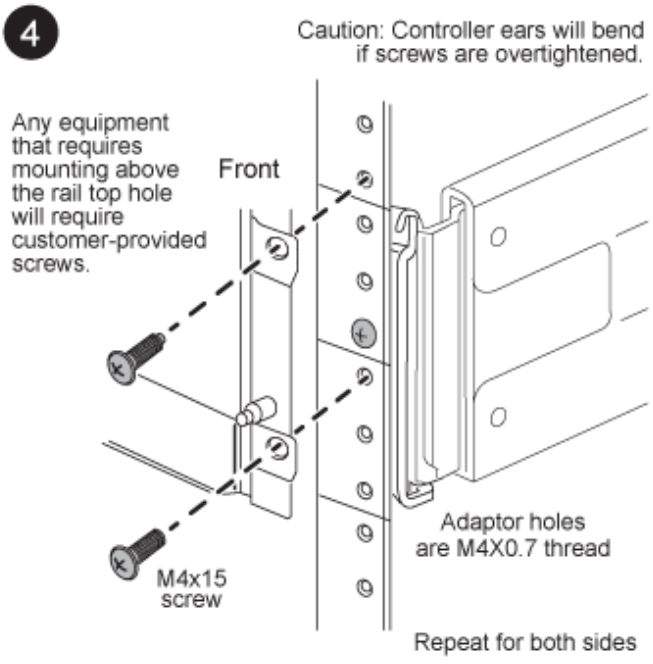

Two-post support rail kit installation instructions - AFF A700 and FAS9000

There are two, two-post support rail kits that can be used with the FAS9000 and AFF A700 systems. One kit allows you to flush-mount your system in the two-post rack, and

the other kit allows you to mid-mount your system in the two-post rack.

Install the two-post mid-mount rail kit

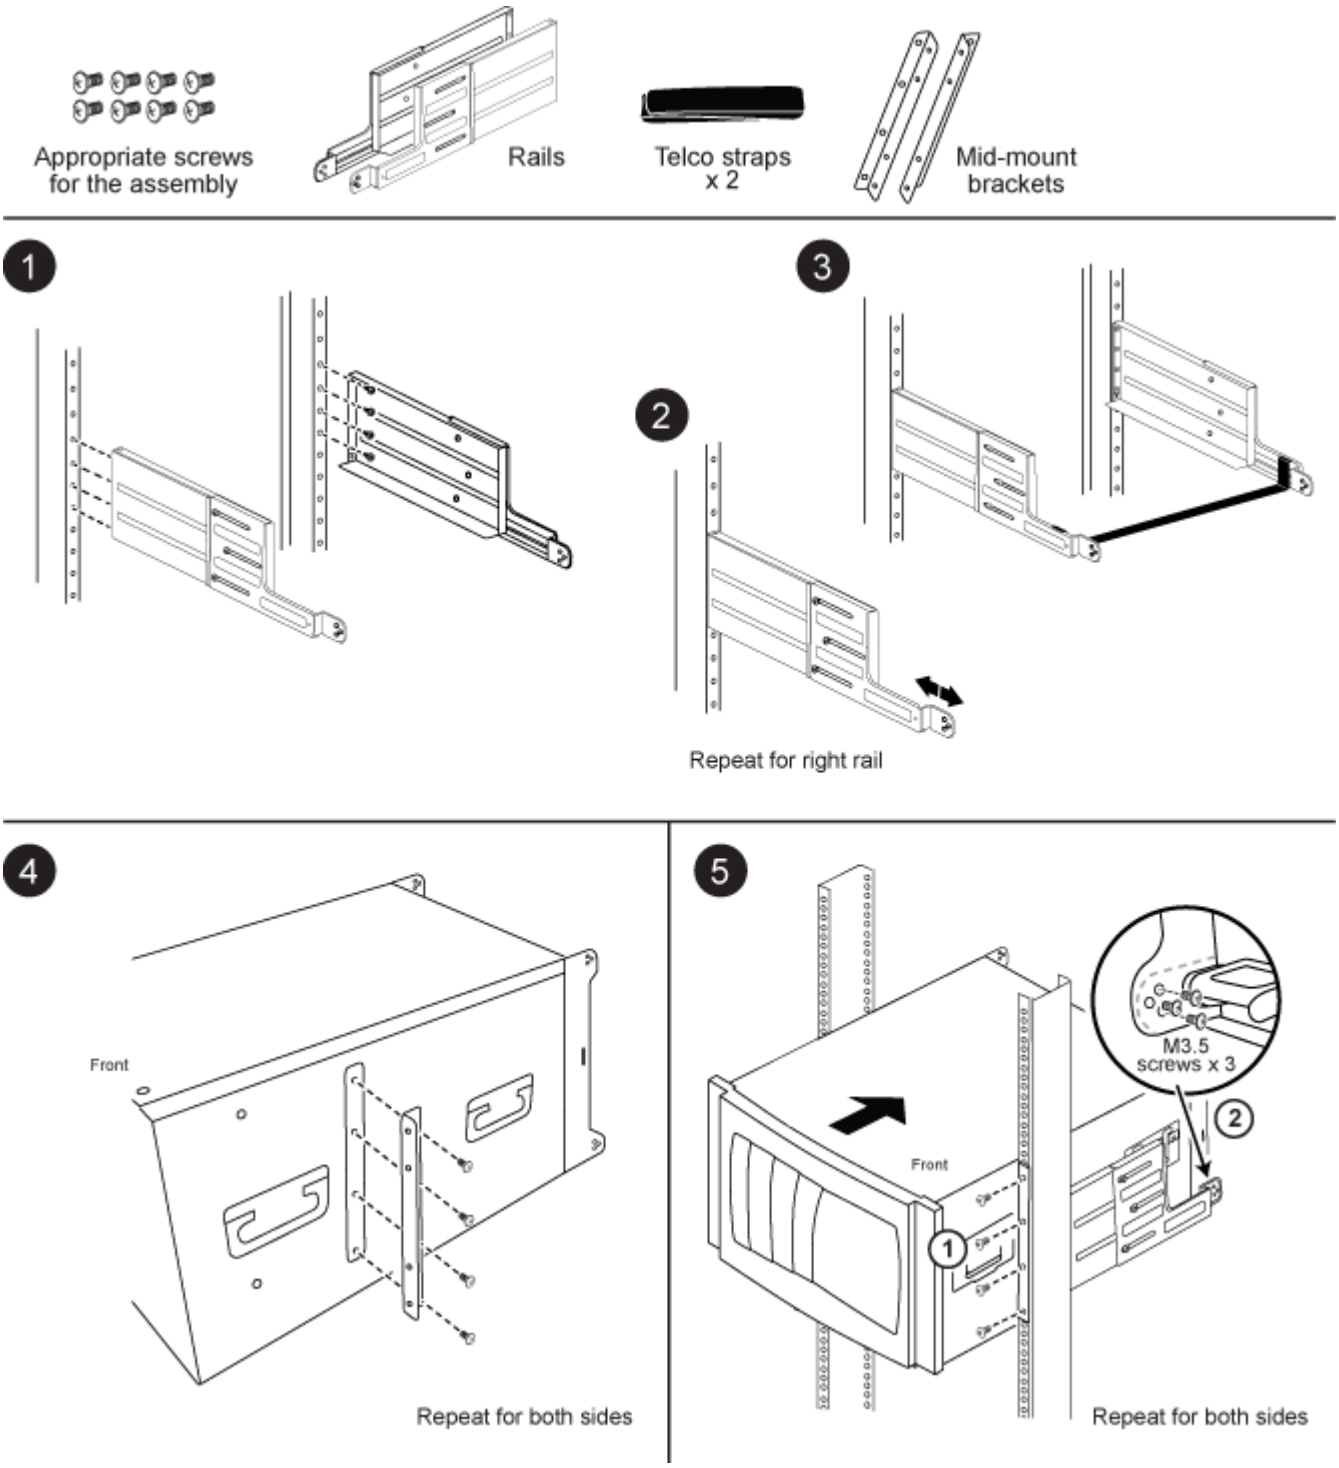

Install the two-post flush-mount rail kit

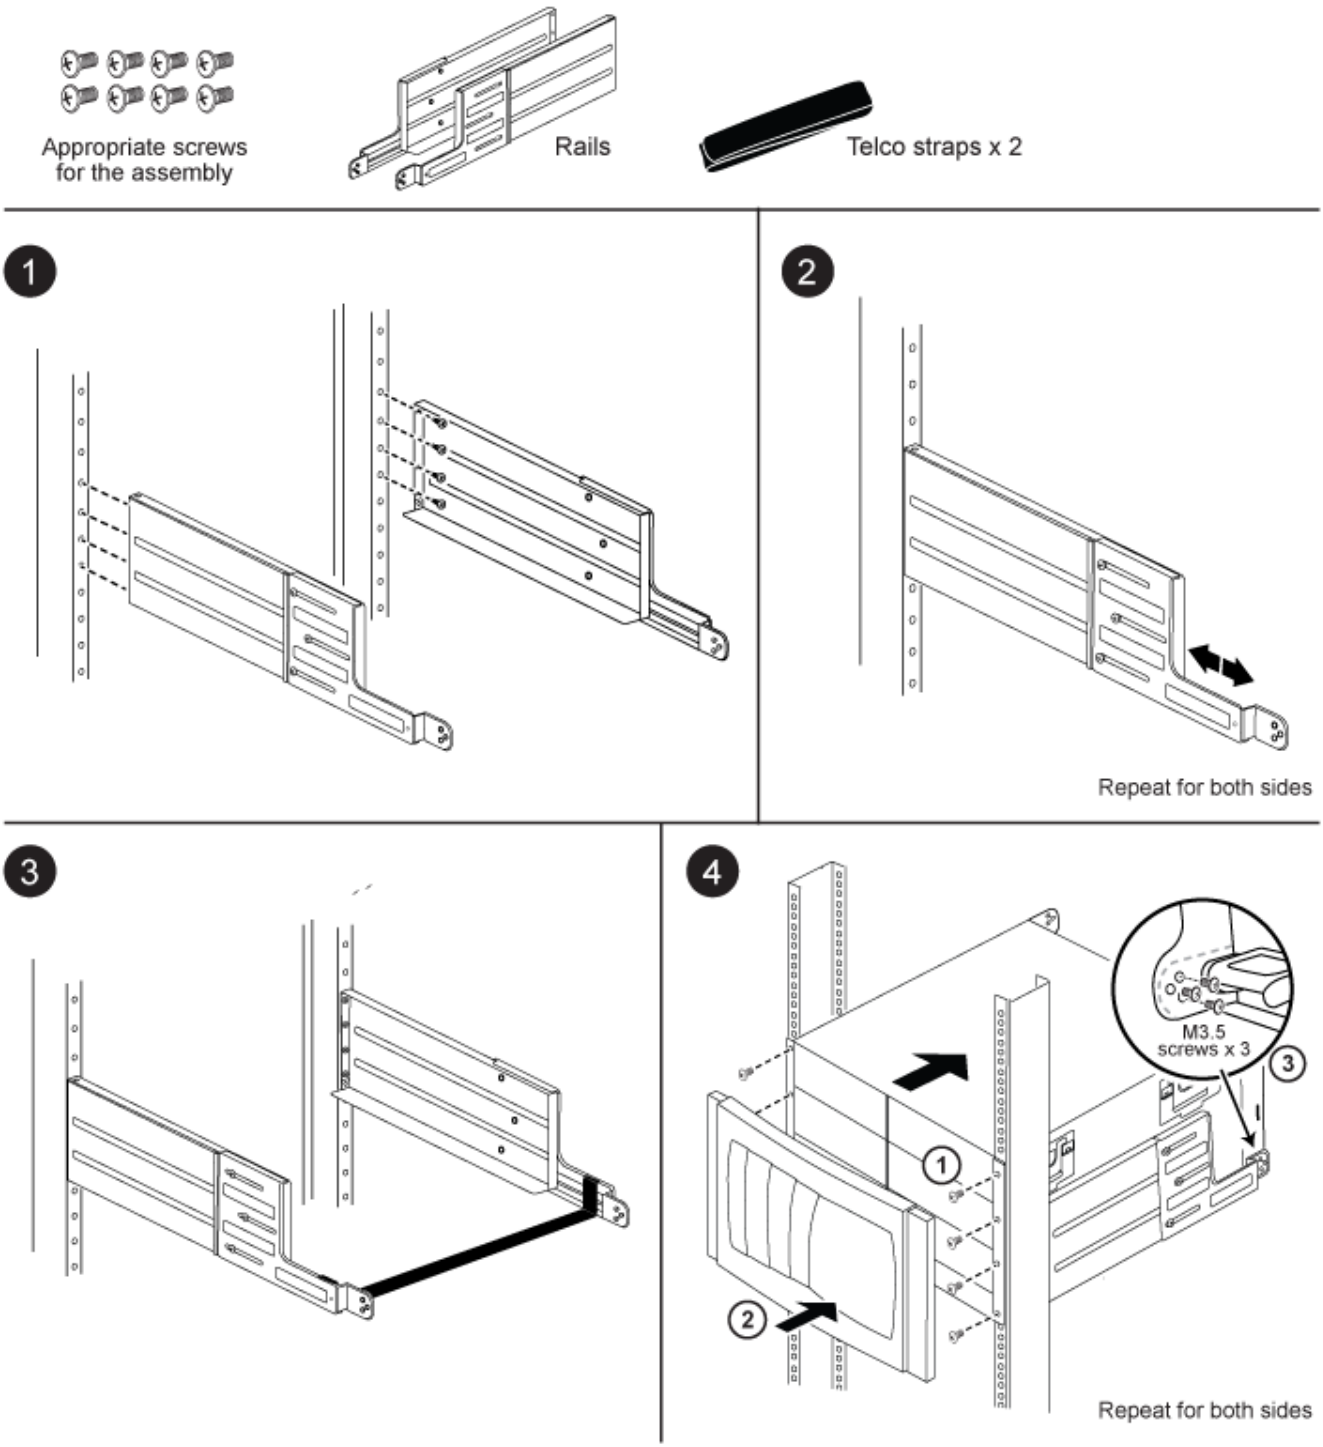

42U 1280 mm system cabinet

Prepare to install cabinet

System cabinet features

The system cabinet consists of side panels, front and rear doors, an optional bolt-down kit, an optional interconnect kit, PDUs for your equipment, and an integrated cable management system.

Feature	Description
Side panels	System cabinets have lockable, removable, and interchangeable side panels.
Perforated front and rear doors	System cabinets have removable front and rear doors with a quick release mechanism. The front door is reversible, and the rear doors are split. Both doors are perforated for cooling.
Common key	This key unlocks the front doors, rear doors, and side panels.
Spares kit	This kit is inside the system cabinet, attached to the cabinet door. It contains the following components: • Four 10-32 x 0.75 inch Phillips pilot screws • Four 10-32 cage nuts • One cage nut insertion tool • Two master key copies
Cable access	Cable pass-throughs are built into the top and bottom of the cabinet, as well as between the bottom of the rear door and the frame.
Cable management	Cable management hook and loop strapping is attached to the frame of the system cabinet at equal intervals.
Support rails	The number of support rails you receive depends on your configuration. The empty system cabinet is shipped with no support rails installed. • For configured system cabinets, one fixed rail kit is shipped with the system cabinet to support the 80xx, FAS8200, and DS4486 rear hold-down brackets. • Quick-ship system cabinets do not include the additional fixed rail kit.
Blanking panels	The number and size of blanking panels you receive depends on your configuration. The empty system cabinet is shipped with no blanking panels installed.
Bolt-down kit	This optional kit enables you to secure the system cabinet to the data center floor. The kit it is not intended for seismic stability. • Four bolt-down brackets • Four spacer brackets • Six M8x20 mm hex head bolts and washers

Feature	Description
Interconnect kit	This optional kit enables you to connect multiple system cabinets to each other. • Interconnect brackets ◦ One set of four interconnect brackets for connecting the system cabinets with side panels on ◦ One set of four interconnect brackets for connecting the system cabinets with the side panels off • Four M12x20 Torx-30 screws used in system cabinet with side panels on. • Eight M6x10 countersunk Torx-30 screws used in system cabinet with side panels off.
Support rail kit	If you ordered additional support rails with your system cabinet, each kit contains one left and one right support rail. The support rails and kit are designed to fit only the NetApp 42U 1280 mm system cabinet. Do not use the rails or a rail kit from other system cabinets because they are not designed for use in the 42U 1280 mm system cabinet. • A left and right support rail • Two screws per rail for securing the rail to the system cabinet frame
Crescent wrench	The crescent wrench is used to remove the hold-down brackets on the packing pallet, adjust the system cabinet leveling feet, and install the bolt-down kit brackets, if ordered.

Required tools and equipment

Before unpacking and installing on your system cabinet, you should gather the necessary tools and equipment to move the system cabinet into place and install it or to perform maintenance on it.

- The appropriate hardware guide for your disk shelves
- The appropriate installation and setup instructions for your system
- #1 and #2 Phillips screwdrivers
- Torq driver for system cabinet screws
- Leveling tool for leveling the system cabinet

Space requirements and system cabinet dimensions

When unpacking your system cabinet, you must make sure that you have enough room to remove the system cabinet from the packing material. Also make sure that the

intended location for the system cabinet is large enough for you to move the cabinet into place.

Required space for unpacking the system cabinet

The following table defines the requires space needed to unpack and install your system cabinet:

Dimensions	U.S.	Metric
Shipping ramp length	80 in.	203.2 cm
Clearance beyond the ramp for cabinet mobility	72 in.	182.9 cm
Shipping pallet depth	59 in.	149.9 cm
Shipping pallet width	42 in.	106.6 cm
Shipping pallet and packaging height	86 in.	218.4 cm
Total rack space, 42U	73.5 in.	186.7 cm
Rail load capacity	Supports all current systems	Supports all current systems
Empty weight	~400 lbs (~181 kg) lbs	~ 181 kg
Fully loaded ship weight	Up to 1,800 lbs	Up to 816.5 kg
Fully loaded static weight	Up to 2,700 lbs	Up to 1,224.7 kg
Front service clearance	47.2 in.	120 cm
Rear service clearance **Note:** The rear door is split. Actual minimum rear clearance is approximately 1/2 the recommendation.	30 in.	76.3 cm
Minimum side clearance for panel removal	24 in.	61 cm
Minimum top clearance	12 in.	30 cm

System cabinet exterior dimensions

The following illustration shows the front, rear, and side views of the system cabinet:

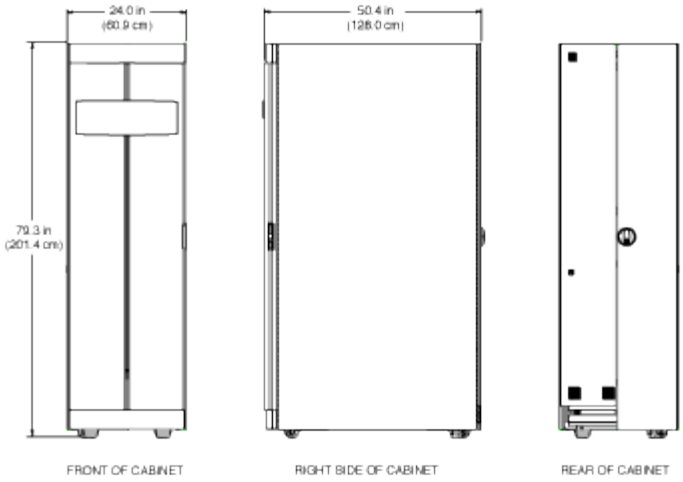

The following illustrations show top and bottom views of the system cabinet, and identify the openings through which you can run cable bundles from the floor of your data center into the system cabinet. The illustrations also show the location of the system cabinet casters and leveling feet.

CAUTION:

To prevent your system cabinet from falling through the data center floor, do not attempt to roll the system cabinet over a floor opening that is wider than the cable access opening at the bottom of the system cabinet.

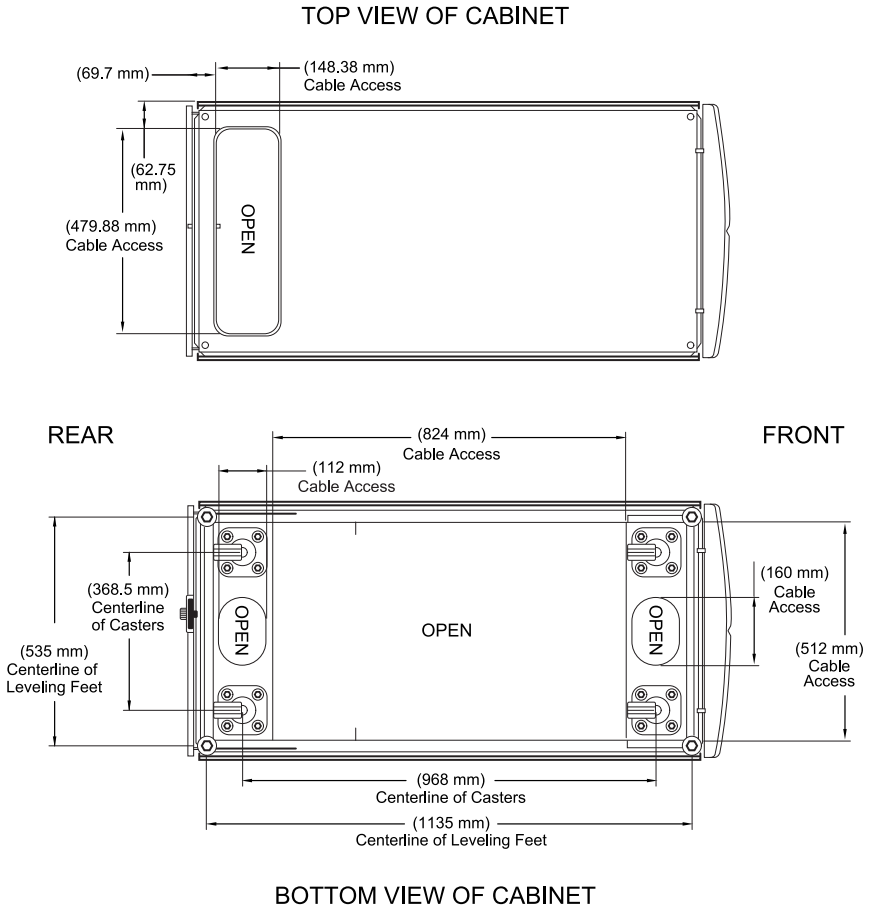

Supported PDU types and specifications

The system cabinet supports different Power Distribution Unit (PDU) types. The PDUs

are compliant with NEMA or IEC.

The most current information for PDUs supported in your system cabinet is listed in the Hardware Universe.

hwu.netapp.com

Unpack the system cabinet

You must remove the packing material that surrounds your system cabinet before you move it into place. You should also recycle the packing material after the cabinet is unpacked.

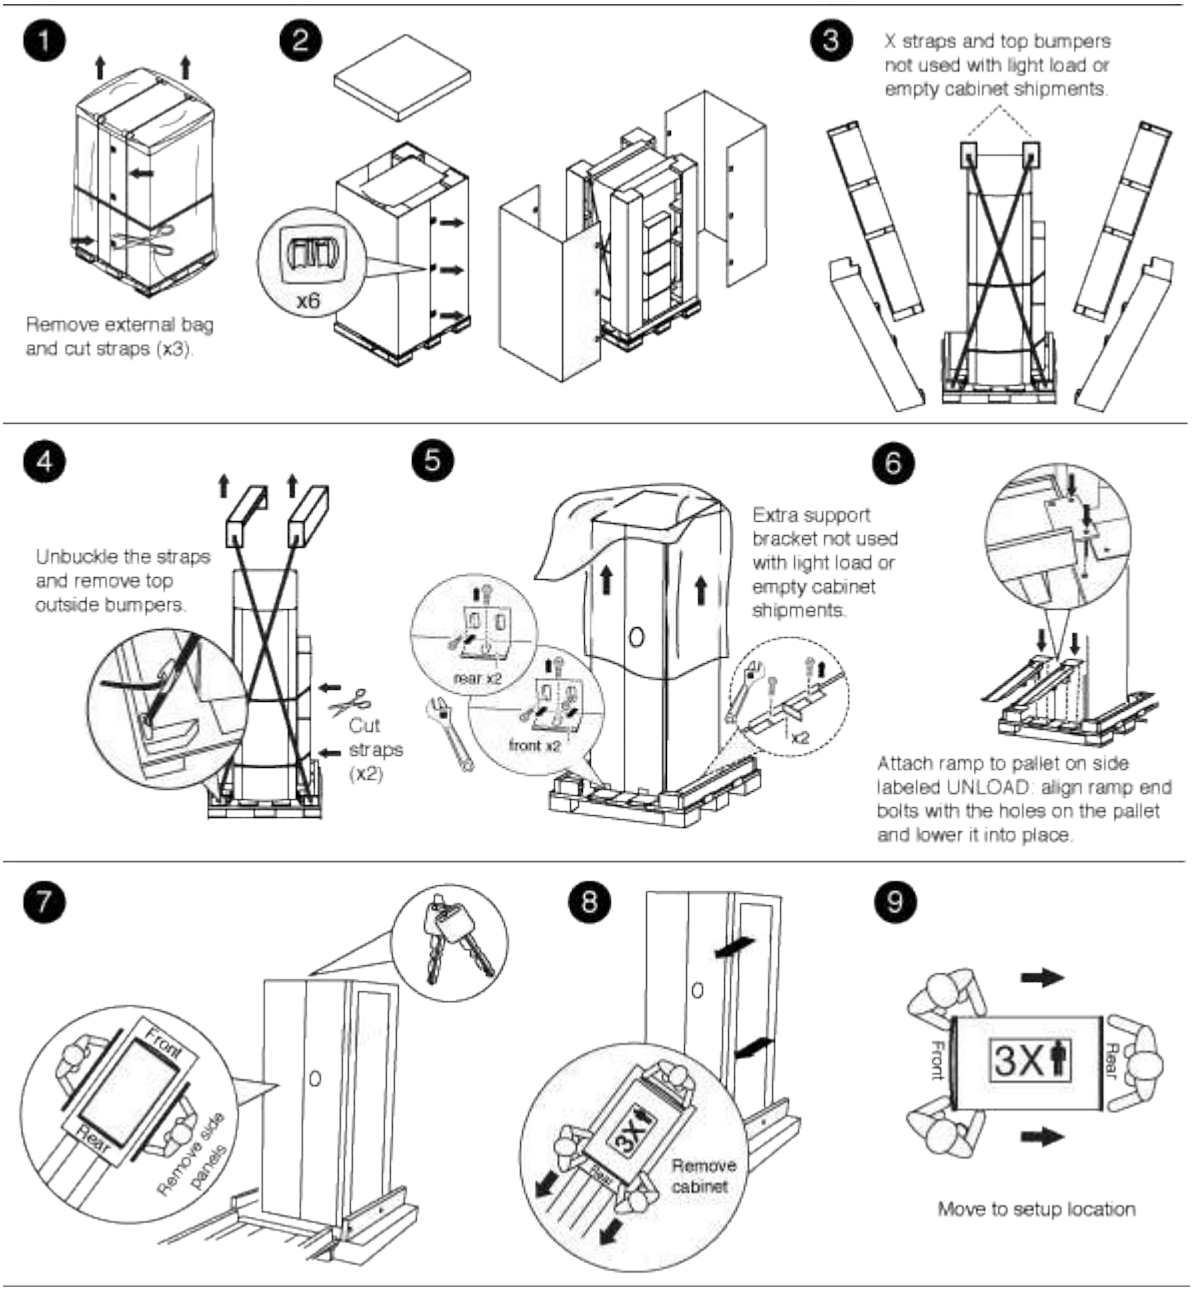

Install cabinet

Install a system cabinet

You can order a system cabinet with NetApp storage controllers and disk shelves installed in it or an empty system cabinet if you already have NetApp equipment. Several system cabinets can be connected together by using the optional interconnect kit, and they can be anchored to the data center floor by using the optional bolt-down kit.

Install the cabinet interconnect kit

You can connect system cabinets together by using the optional cabinet interconnect kit. It is recommended that you install the kit to prevent the cabinets from pulling apart and damaging system cables.

1. Place the system cabinets close together.

   The cabinets should be arranged similarly to the following illustration, with the cabinet with the controller modules in the middle, and the cabinets with additional disk shelves on either side. The sides of the cabinets should be close, but do not need to touch each other yet.

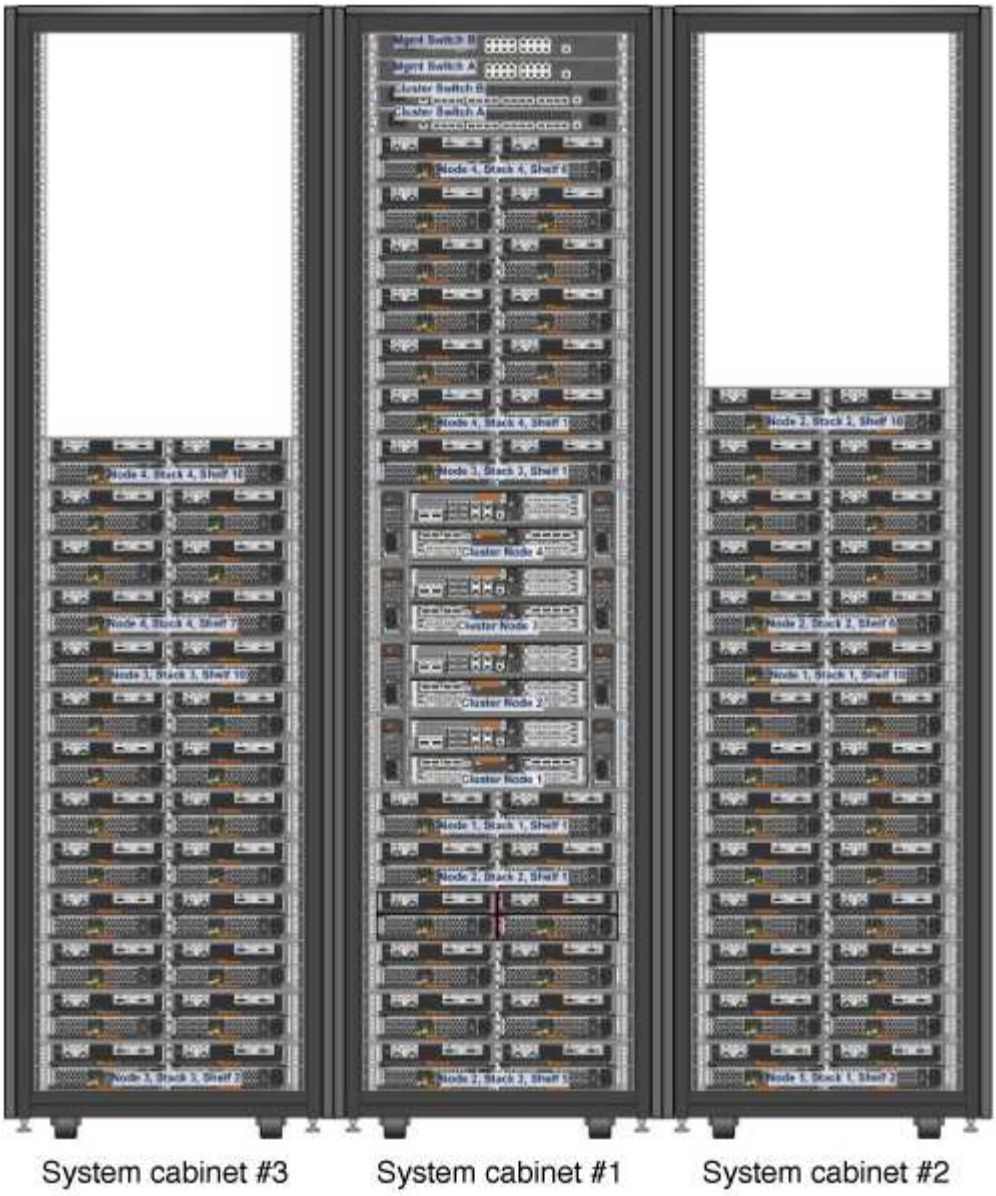

2. If you are installing the interconnect kit with the side panels on as recommended, reinstall the side panels that were removed during unpacking:
   a. Lift the side panel, tilting it about 15 degrees away from the system cabinet bottom, and then hang it over the lip at the top of the system cabinet frame.
   b. Gently push the side panel against the cabinet frame, and then lock it in place with the key.
   c. Repeat these substeps for the remaining side panels.
3. If you are installing the interconnect kit with the side panels removed, remove the front door whose hinges are on the edge where the cabinets meet:
   a. Unlock and open the front door that is being removed.
   b. Use the following illustration for reference to unplug the power to the illuminated bezel:

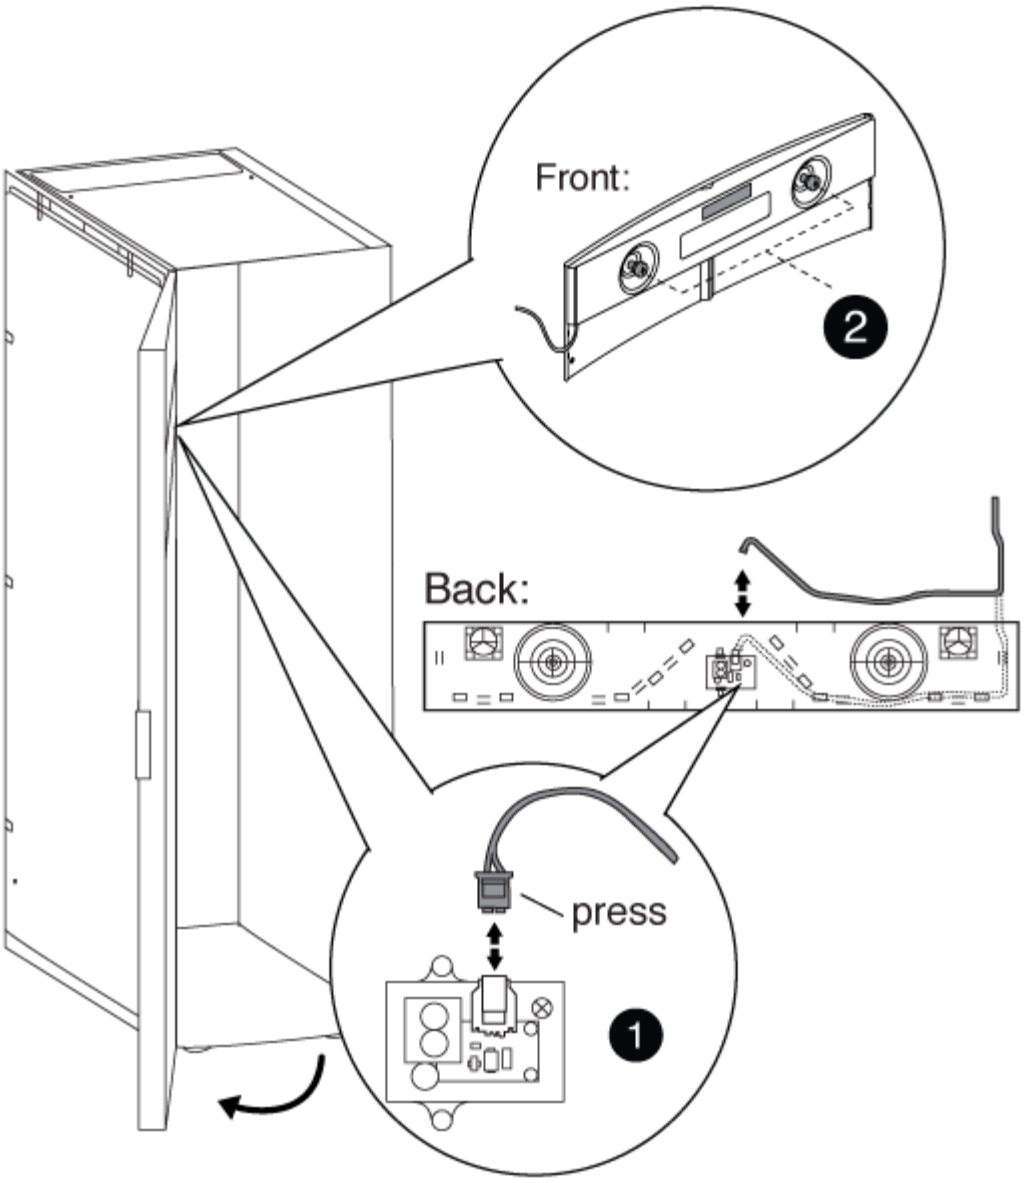

1
Illuminated bezel circuit board and cable
2
Back panel and thumbscrews

c. Use the following illustration for reference to remove the front door:

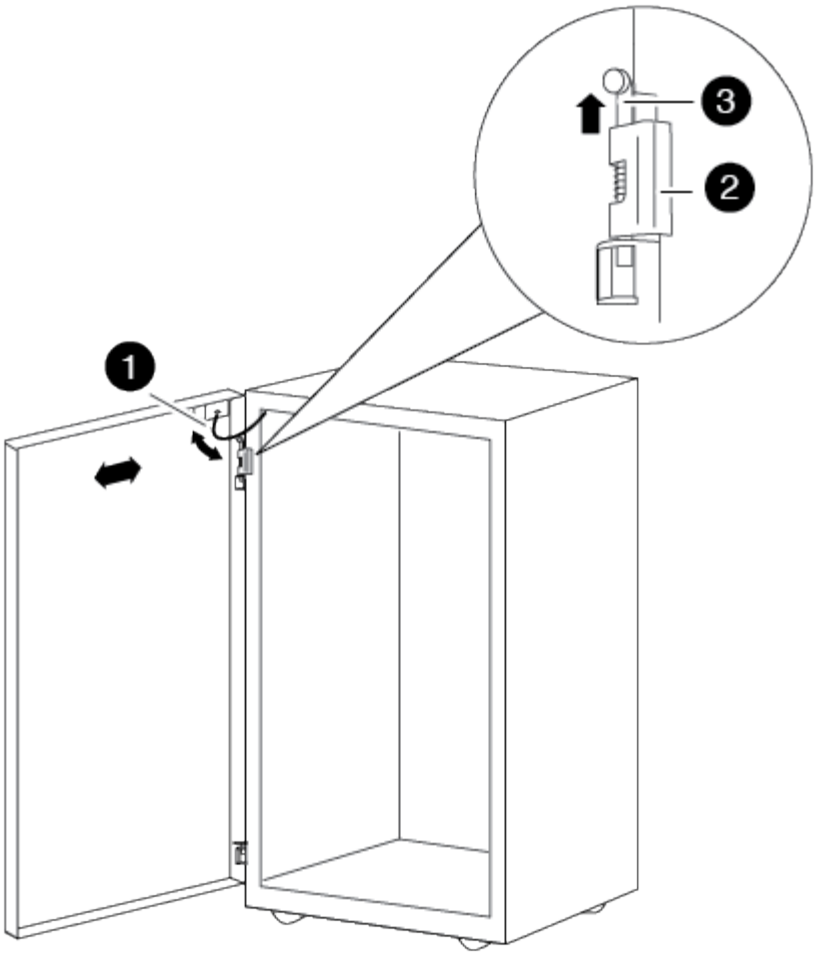

1	
Door grounding cable	
2	
Door top hinge	
3	
Hinge pin	

Make sure that you set the removed doors in a safe place so that they are not accidently damaged.

4. Remove the rear door whose hinges are on the edge where the cabinets meet:
   a. Unlock and open the rear door that you are removing.
   b. Lift the top hinge pin until it clears the bottom of the hinge.
   c. Gently tip the top of the door away from the system cabinet frame, and then release the hinge pin.
   d. Lift the door off the bottom hinge, and then set the door aside.

5. Move the system cabinets completely together, and then align and level them by adjusting the four leveling feet at the bottom of the system cabinets.
6. Install the interconnect brackets.
   - Use the following illustration for reference if you are installing the interconnect brackets with the system cabinet side panels on, as recommended:

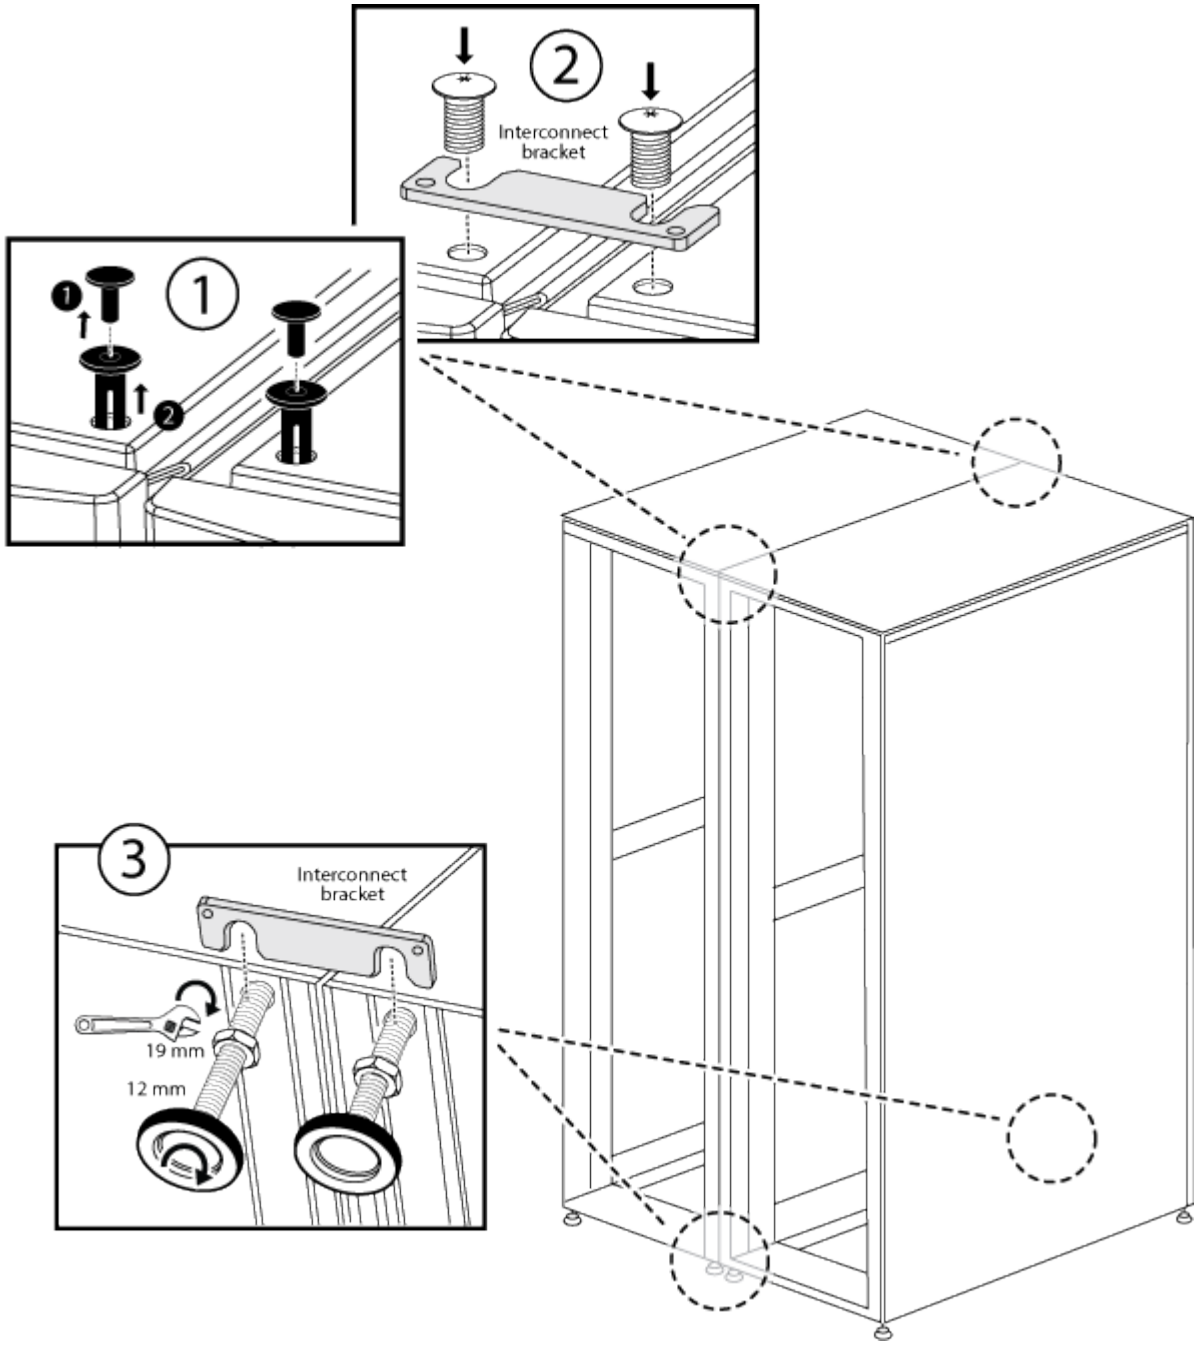

1	
Plastic push-in rivets on the system cabinet top	
2	

Top interconnect bracket	
3	
Bottom interconnect bracket	

- Use the following illustration for reference if you are installing the interconnect brackets with the system cabinet side panels off:

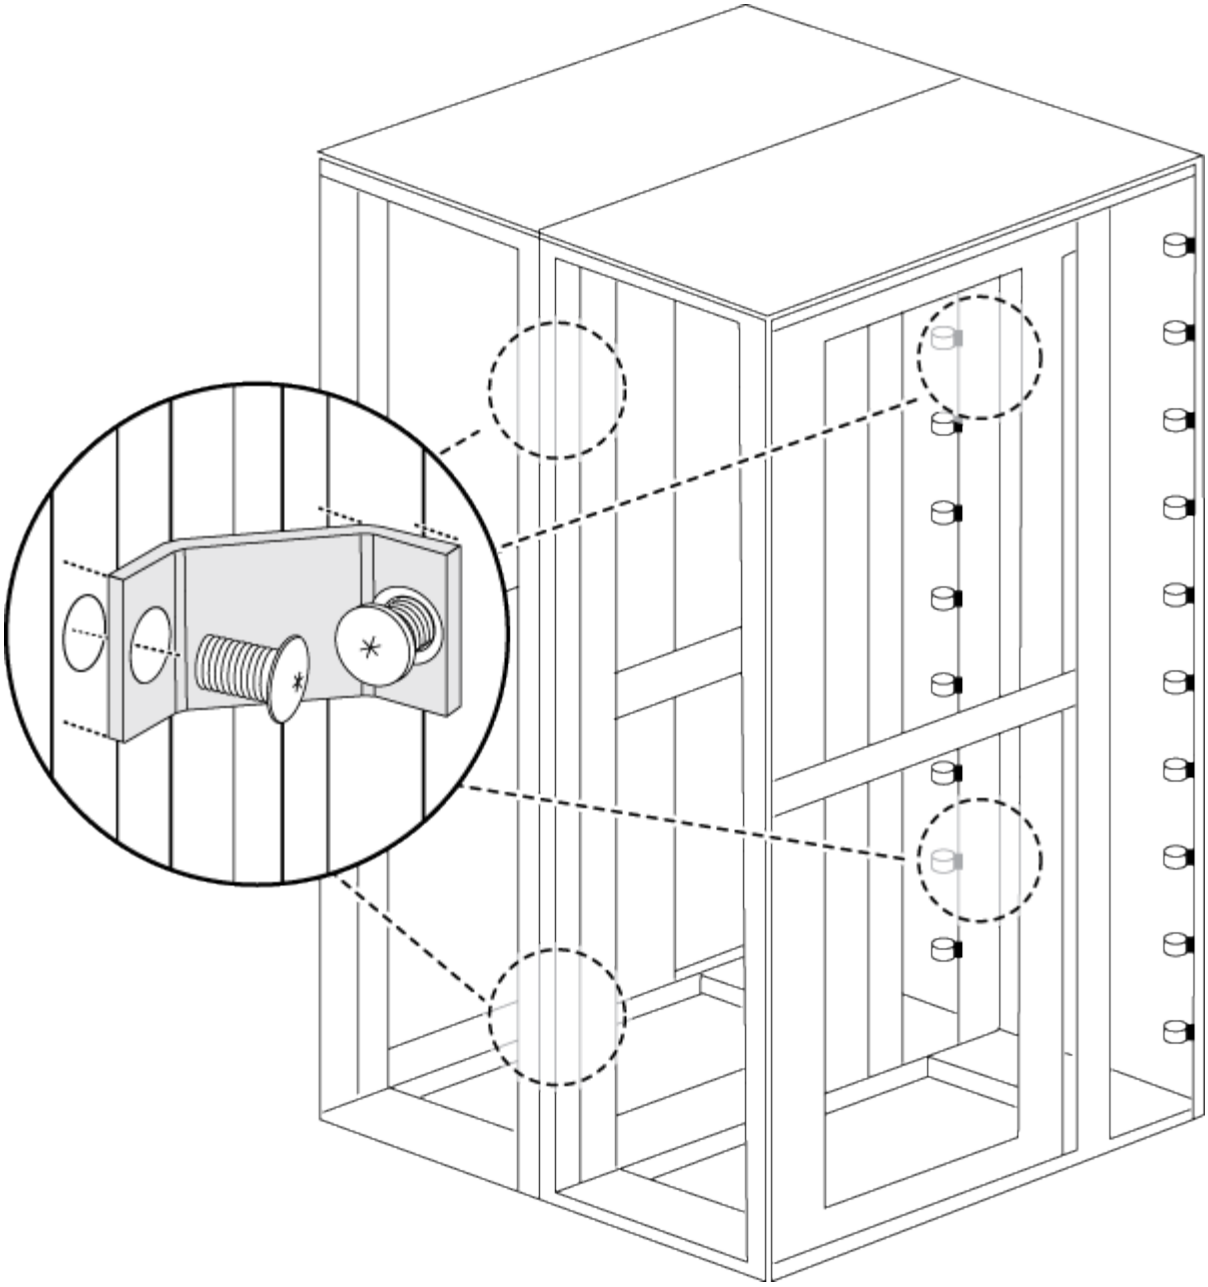

7. Repeat the process for any remaining system cabinets.
8. Tighten all interconnect bracket screws.

Install the bolt-down kit

You can secure the system cabinet to the floor by installing the optional bolt-down kit. Installing the kit prevents the system cabinets from being rolled out of position.

You must supply the appropriate anchor bolt for your floor for each bolt-down bracket.

1. Mark the area on your floor where the system cabinet will be installed, and then roll the cabinet into place.

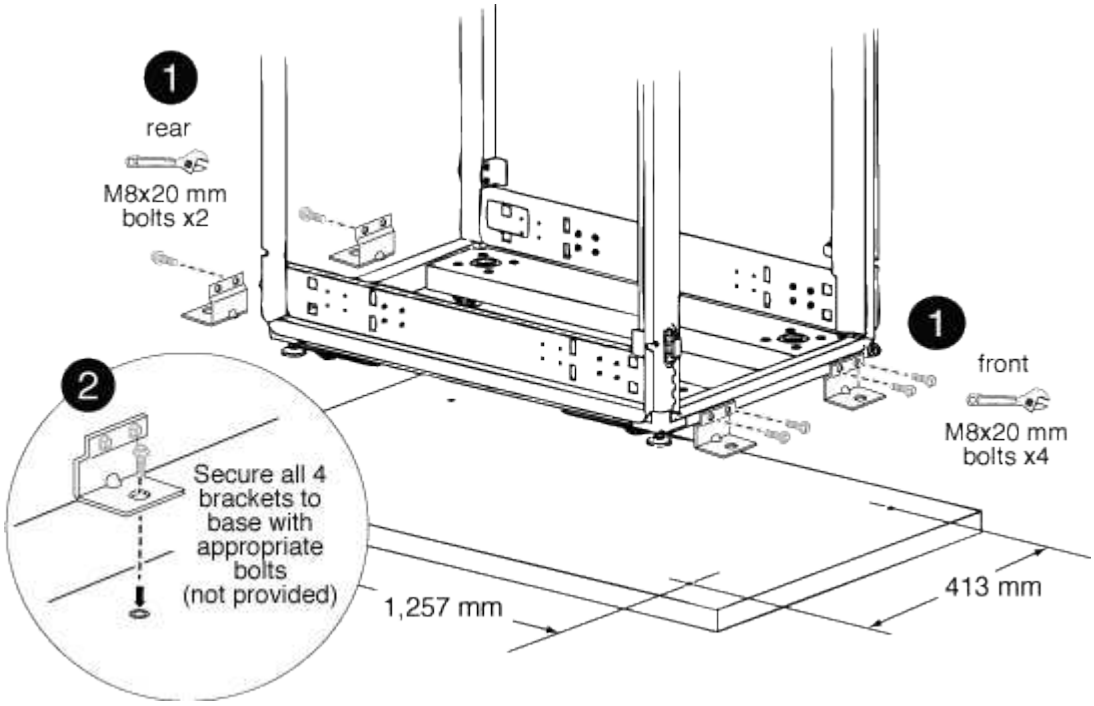

1
Front and rear bolt-down brackets
2
Location of floor anchor point on the bracket

2. Mark the anchoring points where the rear bolt-down brackets will be anchored to the floor, and then drill the holes for the brackets.

   Be sure to use the appropriate bolt sizes and type for your floor.

3. If the bolt-down brackets are too low to align with the mount points on the system cabinet frame, place a spacer bracket over the hole in the floor.
4. Loosely bolt the rear brackets to the floor, and then using the kit bolts, bolt the brackets to the cabinet frame.
5. Mark the anchoring points where the front bolt-down brackets will be anchored to the floor, and then drill the holes for the brackets.

6. If the bolt-down brackets are too low to align with the mount points on the system cabinet frame, place a spacer bracket over the hole in the floor.
7. Bolt the front brackets to the floor, and then using the kit bolts, bolt the brackets to the cabinet frame.
8. Lower the leveling feet as needed, and then tighten the rear bolt-down brackets to the floor.

Install additional support rails

Your system cabinet has some support rails already installed in it. If you need additional support rails for your system, you must install them before installing your system components.

This task applies to all controller and disk shelves except the DS212C and the DE212C disk shelves. Use the instructions in the rail kit flyer applicable to those two disk shelves.

Installing a DE212C or DS212C Shelf in a Two-Post or Four-Post Rack

1. Determine how much space your equipment requires.

   Calculate the amount of U space (1.75 inches per U) the equipment requires, based on the equipment height, and then determine where the equipment will be installed in the system cabinet based on available space.

2. Locate where you need to install the support rails, and then install them using the following illustration for reference:

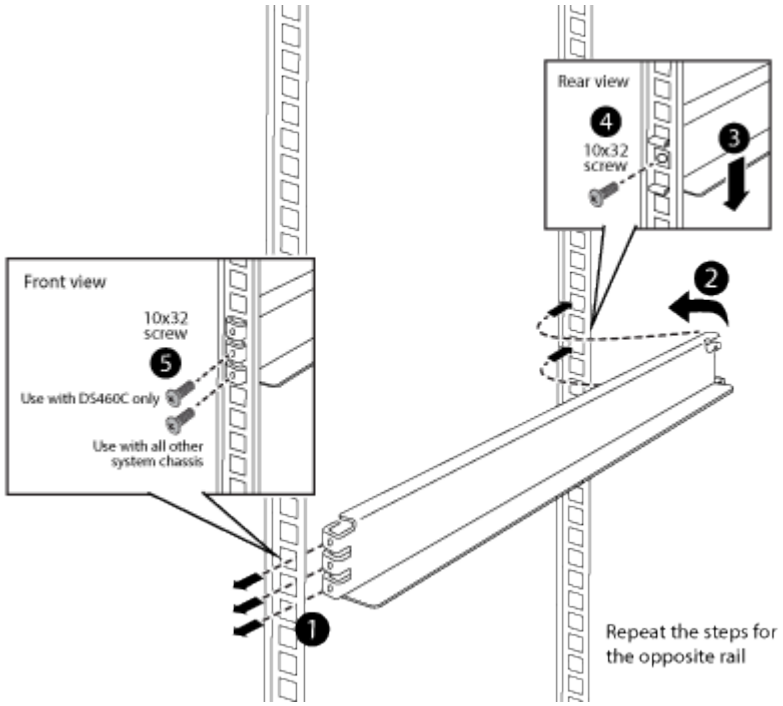

3. If your equipment mounting flanges extend beyond the screw holes in the support rail, install cage nuts above the support rail, where needed.

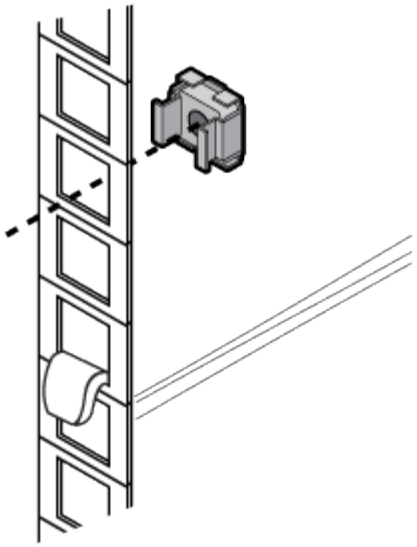

Install equipment in the system cabinet

After you have installed any additional support rails into the system cabinet, you can add more system components to your prepopulated system cabinet or add your existing system components to an empty system cabinet.

1. Unlock and open the rear doors of the system cabinet and the front door, if it is not already open.
2. Install your equipment into the system cabinet as described in the installation instructions accompanying your equipment.

   The storage controllers should be in the middle of the system cabinet. The disk shelves should be above and below the storage controllers. Any switches should be at the very top of the system cabinet.

   If your equipment mounting flanges extend beyond the screw holes in the support rail, install cage nuts above the support rail where needed to secure the equipment to the cabinet upright.

3. Install blanking panels over any empty bays in the system cabinet.

   If you receive the system cabinet with equipment already installed, you must remove the tie-down rails on top of the equipment that is directly below empty cabinet bays, as shown in the following illustration:

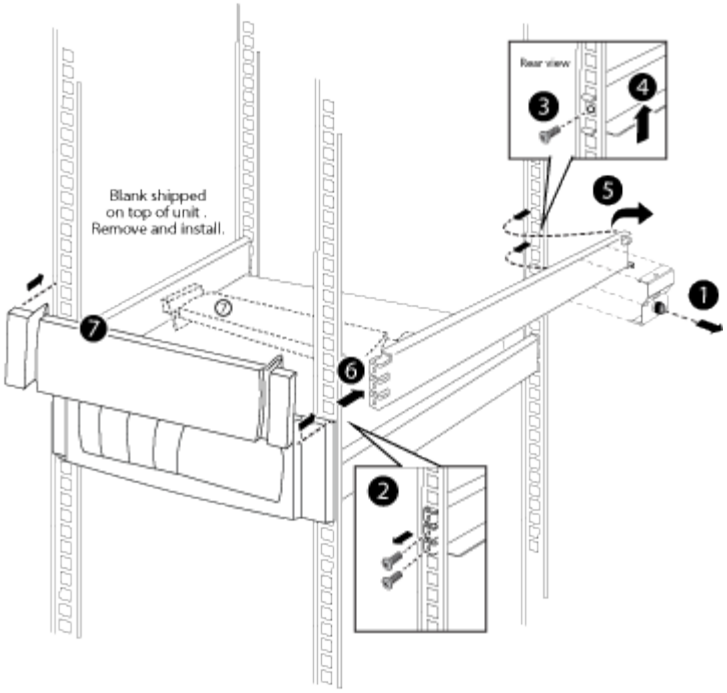

4. Reinstall the front and rear system cabinet doors.

Power on the system cabinet

You must connect the system components to the PDUs, route the PDU cables to the AC power sources, connect them to the power sources, and power on the system.

To prevent a system outage if a PDU fails, do not plug both power cables from any component into one PDU. Distribute the load across the PDS that are plugged into different circuits, as shown in the illustration.

You must have separate power circuits available for each PDU in your system cabinet.

1. Connect your equipment to the PDUs:
   a. Connect the controller power supplies to different PDUs connected to different power sources.
   b. Connect the drive shelf power supplies to different PDUs connected to different power sources.
   c. Connect any switch power supplies to different PDUs connected to different power sources.

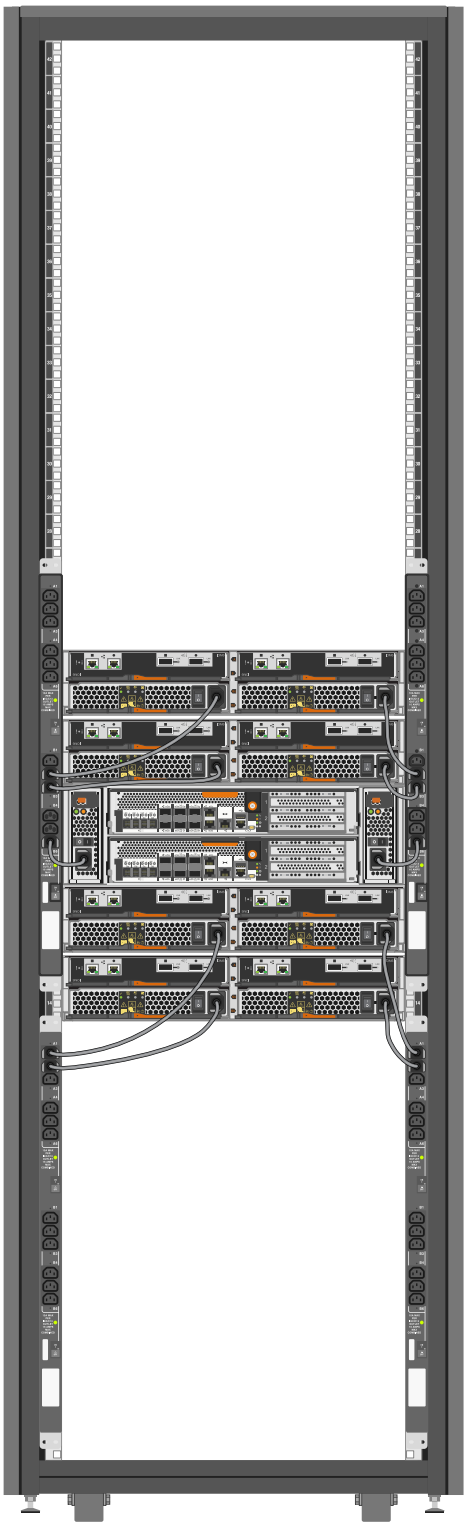

Feed the PDU power cables through one of the following openings in the system cabinet:

- The top of the system cabinet
- Between the rear door bottom and frame of the system cabinet
- Through the floor opening and under the system cabinet

2. Turn off the power switches or circuit breakers on the PDUs.
3. Plug each PDU power cable into individual AC power sources that are on separate AC circuits.

4. Turn on the power switches or circuit breakers to the PDUs.
5. Turn on the power to your components, and then boot the system.
6. Close and lock the system cabinet doors.

Replace PDUs

You can replace a failed PDU in your system cabinet or replace an existing PDU with a different type of PDU.

The replacement PDU must be supported by your system cabinet and must provide sufficient power to the installed equipment.

hwu.netapp.com

1. Turn off the circuit breakers on the impaired PDU, and then unplug the impaired PDU from the AC power source.
2. Ground yourself to the system cabinet, and then unplug the power cords from each of the system components and from the impaired PDU.
3. Remove the screws from the PDU frame, bottom screw first.

   Ensure that you support the PDU with one hand while you remove the last screw from the top of the PDU. This prevents the PDU from dropping or falling toward you after the screw is removed.

4. Remove the impaired PDU from the system cabinet.

   Make sure that you keep track of the mounting screws so that you can reuse them when installing the replacement PDU.

5. Remove the brackets from the impaired PDU, and then install them on the back of the replacement PDU.
6. While supporting the replacement PDU, align the slot on the mounting bracket of the PDU with the top holes of the frame on the inside of the system cabinet, and then secure the PDU to the system cabinet frame using the mounting screws from the impaired PDU.
7. Secure the bottom of the PDU to the system cabinet frame, and then tighten all of the mounting screws.
8. Verify that all of the power switches or circuit breakers are in the Off position.

   If the circuit breakers are not in the Off position, push a small screwdriver or straightened paper clip into the slot to the right of the Off label to trip the circuit breaker and turn off the circuit.

9. Plug the power cords you unplugged from the storage system, switches, and drive shelves into the replacement PDU, plugging each component into the PDU outlet directly across from the component.

   A best practice is to distribute the total load across the PDU branches, making each branch load as equal as possible.

10. Lock each component power cable plug in place with the cable retainer clip above it by sliding the curved edge of the cable retainer clip over the plug shoulder.
11. Plug the PDU power cord into the AC power source.
12. Turn on the PDU power switches or PDU circuit breakers.

For PDU circuit breakers, the button is on when it is flush with the PDU frame.

Reverse cabinet front door

Reverse the system cabinet front door

You can change the direction the front door opens by removing the illuminated badge, door, top hinge, and related hardware, and then installing them on the opposite side of the front of the system cabinet frame.

You need the following tools and equipment to complete the door reversal for system cabinets with illuminated badges:

- A Phillips screwdriver
- A 5-mm Allen wrench; magnetic Allen wrench is recommended
- Needle-nose pliers
- A step ladder so that you can easily access the Allen bolts in the top hinge

Remove the illuminated badge

Removing the illuminated badge requires that you open the system cabinet front door, unplug the power cord from the back of the badge, and then remove the badge components from the system cabinet door.

Use the following illustration along with the following steps:

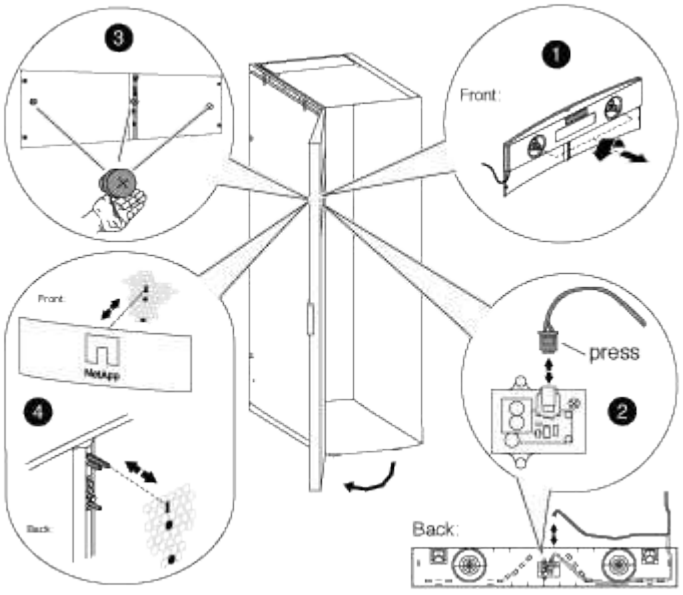

1. Unlock and open the system cabinet front door.
2. Loosen the captive screws on the badge back panel on the inside of the door, and then gently pull the back panel away from the door mesh.
3. Unplug the power cord from the back panel by pressing the locking clip on the plug, unplugging the cord from the socket, and removing the cable from the back panel.

   Set the back panel aside.

4. Carefully remove the screws from the back of the badge.

> The stems on the thumbscrews are very short. Place your free hand under the screw to catch the thumbscrew if you drop it.

5. Remove the badge from the front of the door and set it aside.

Remove the system cabinet door

You must remove the system cabinet door and side panels to move the illuminated badge and components, and to reverse the door.

1. Open the system cabinet door if it is not already open.
2. Perform the appropriate action depending on whether your cabinets are connected with the interconnect kit.

If your system cabinet is…	Then…
Not connected to another system cabinet	Go to the next step.
Connected to another system cabinet with an interconnect kit	Remove all four interconnect kit brackets and set the brackets and screws in a safe place.

3. Unlock both side panels, disconnect the grounding wires from the side panels, and then remove them and set them aside.
4. Disconnect the grounding wire from the grounding spade located at the top of the door.
5. Unscrew the grounding lug and wire assembly from the system cabinet frame and set it aside.
6. Unscrew the grounding lug assembly from the system cabinet door and set it aside.
7. Lift the top hinge pin until it clears the bottom of the hinge.
8. Gently tip the top of the door away from the system cabinet frame, and then release the hinge pin.
9. Lift the door off the bottom hinge, and set the door aside.

Move the badge power supply and cabling

You must move the power supply and illuminated badge cabling to the opposite side of the system cabinet frame before you reverse the door and reinstall the illuminated badge.

You must have removed the system cabinet door and side panels.

You must move the illuminated badge power supply, power cable, and cabling conduit to the opposite side of the system cabinet when you reverse the system cabinet door. The assembly is designed so that the cable to the badge is on the side of the cabinet where the door hinge is installed.

1. Open the power cable retaining clip, and then disconnect the power cable from the power supply.
2. Remove the power supply housing and power supply, using the illustration for reference:

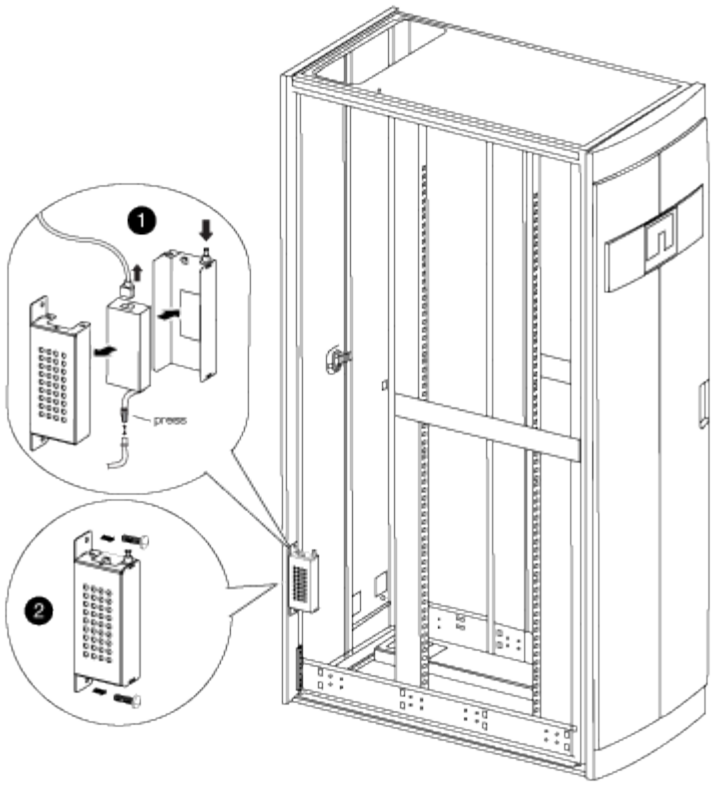

a. Lift the retaining pin on the power supply housing, and then remove the housing cover by rotating it downward and lifting it off the rear power supply housing.

> The power supply is attached to the power supply housing with a hook and loop patch.

b. Disconnect the power supply from the illuminated badge cable, and then set the power supply and power supply cover to the side.

c. Remove the screws from the top and bottom of the power supply housing that is attached to the system cabinet frame, and then remove the power supply housing.

3. Install the power supply and power supply housing on the opposite side of the system cabinet:

a. Locate the two screw holes next to each other on the cabinet frame, and then attach the top of the power supply housing to the bottom-most of the two screw holes.

> You might need to remove the bottom cable retention strap, if present.

b. Secure the bottom of the power supply housing to the system cabinet frame.

c. Install the power supply cover and power supply by aligning the cover hooks with the power supply back, pulling the plunger up on the cover, rotating the plunger closed, and then releasing the plunger.

4. Remove the bezel power supply conduit by removing the conduit retaining clips from the retaining clips, and then slide the conduit off the power cable.

Keep the retaining clips and screws for installing the conduit on the opposite side of the cabinet.

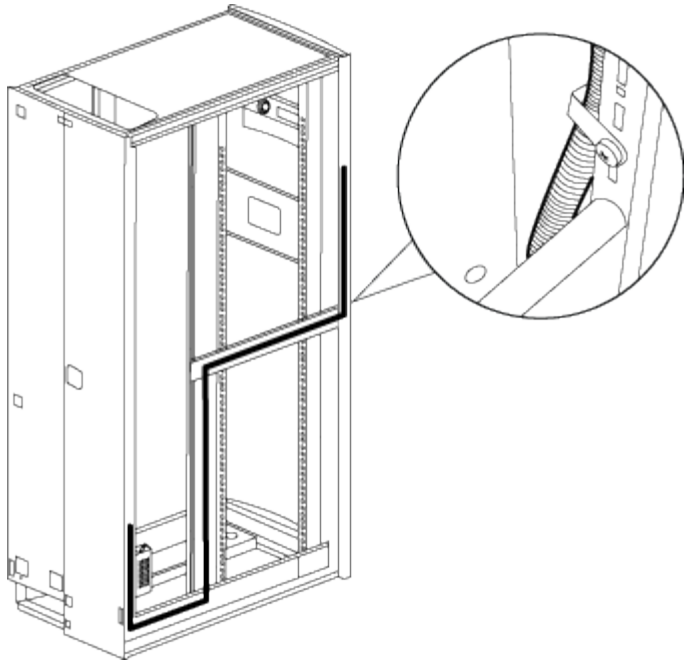

5. Move the badge power cable to the other side of the cabinet:

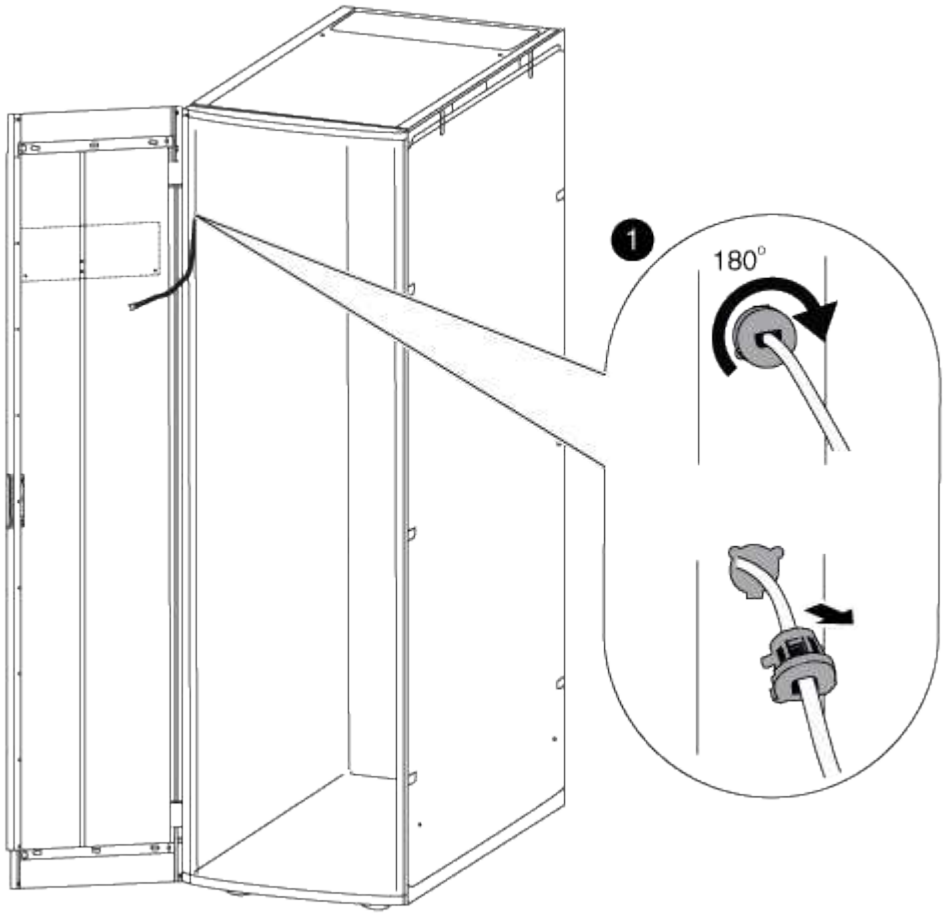

a. Rotate the rubber cable retainer on the cabinet upright 180° to the right, remove it from the system cabinet frame, and then gently pull the cable out of the system cabinet.

b. Move the cable to the other side of the cabinet, and then thread it completely through the hole near the top of the cabinet upright.

c. Align the rubber cable retainer with the hole in the frame, push it in as far as it will go, and then rotate the cable retainer 180° to the left to secure it.

d. Run the cable along the cabinet frame to the back of the cabinet.

6. Reinstall the cable conduit:

a. Slide the conduit over the PDU power cable and route the conduit along the system cabinet frame to the PDU.

b. Install the conduit retaining clips from the other side of the cabinet over the conduit to secure it to the cabinet frame.

7. Plug the badge cable back into the power supply, but do not reconnect the power supply to the power source.

Reverse the door hinge and lock catch

When reversing the system cabinet door, you must move the system cabinet door hinge and lock catch to the opposite front-side system cabinet upright.

You need the following tools:

- Phillips screwdriver
- 5 mm Allen wrench; magnetic Allen wrench is recommended
- Needle-nose pliers
- Step ladder so that you can easily access the Allen screws in the top hinge

1. Remove the screws securing the top hinge from the system cabinet frame, and set the screws and hinge aside.

> Be careful when removing the Allen screws to avoid dropping them into the cabinet frame. Spare Allen screws are provided in the spares kit that shipped with your system cabinet.

2. Remove the screws securing the bottom hinge from the system cabinet frame, and set the screws and hinge aside.

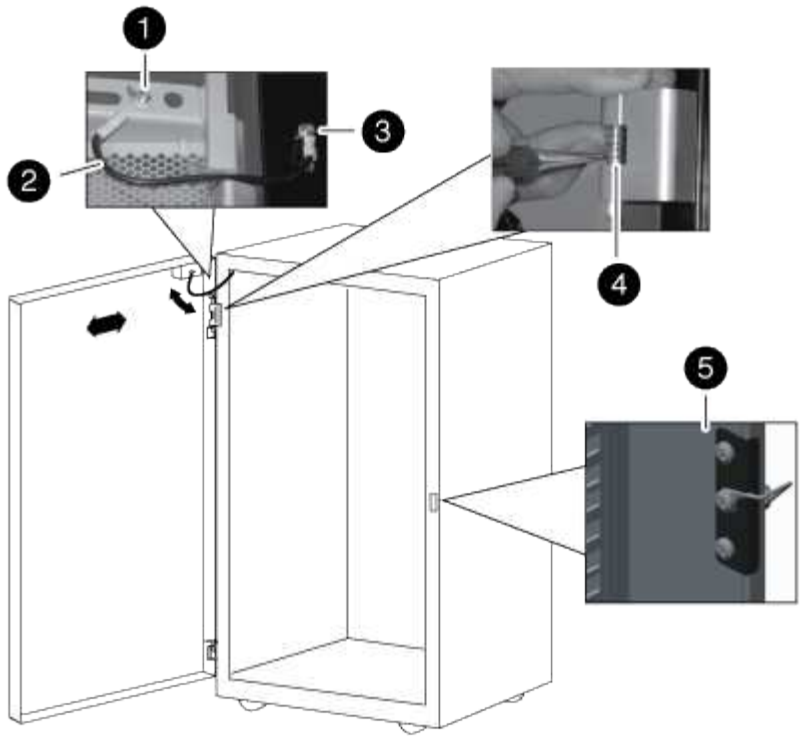

1
Door grounding screw with grounding wire spade
2
Grounding wire
3
Frame grounding wire lug
4
Top front door hinge with hinge pin held by retaining clip
5
Lock catch

3. Reverse the hinge pin from the top hinge:
   a. Lift the hinge pin and expose the retaining clip on the hinge pin shaft.

b. Using the needle-nose pliers, gently remove the retaining clip from the hinge pin shaft and set it aside.

c. Slide the hinge pin and spring out of the hinge body.

d. Rotate the hinge so that the thread holes are facing the opposite side of the hinge, and then install the hinge pin and spring back into the hinge.

e. Install the hinge retaining clip onto the hinge pin.

   Make sure that you push the retaining clip completely onto the hinge pin.

4. Reinstall the hinges:

   a. Insert the top Allen screw through the system cabinet upright, aligning it with the top threaded hole on the top hinge, and then partially tighten the Allen screw.

      Do not completely tighten the screw until after the second Allen screw is installed.

   b. Insert the bottom Allen screw through the system cabinet upright, aligning it with the bottom threaded hole on the top hinge, and then partially tighten the Allen screw.

   c. Tighten the top and bottom Allen screws.

   d. Repeat these steps for the bottom hinge.

5. Remove the screws from the lock catch, and then move the lock catch to the opposite front-side system cabinet upright.

6. Rotate the catch 180 degrees, and then secure it to the system cabinet upright.

Reinstall the door and illuminated badge

After you move the power supply and components to the other side of the system cabinet and moved the hinges and lock catch, you must reinstall the system cabinet door and the illuminated badge, and then reconnect the badge to the power source.

Reinstall the system cabinet door

After you reverse the door hinge and door catch, you must reinstall the grounding wire and lug assembly and wire, and the system cabinet front door prior to reinstalling the illuminated badge.

1. Rotate the door 180 degrees.
2. Align the bottom of the door with the bottom hinge post, and then seat the door bottom on the hinge post.
3. Lift the top hinge pin so that it clears the hinge housing.
4. Tip the top of the door into the hinge housing so that the hinge pin and door hinge are aligned, and then release the hinge pin.

   Make sure that the hinge pin is seated completely through the door hinge and the bottom of the door hinge housing.

5. Reattach the grounding lug and wire assembly to the system cabinet frame on the same side of the newly reversed front door and reinstall the grounding lug with spade on the top of the system cabinet door.
6. Reattach the grounding wire to the spade on the grounding lug assembly on the system cabinet door.
7. Reinstall either the side panels or the interconnect brackets, as applicable:

- If your system cabinet is not connected to another system cabinet, reinstall the side panels.
- If your system cabinet is connected to another system cabinet with an interconnect kit, reinstall the interconnect brackets.

Reinstall the illuminated badge

After the system cabinet door is installed, you need to install the illuminated badge to complete the door reversal process, and then close and lock the front door.

1. Using the following illustration for reference, reinstall the illuminated badge on the front door of the system cabinet:

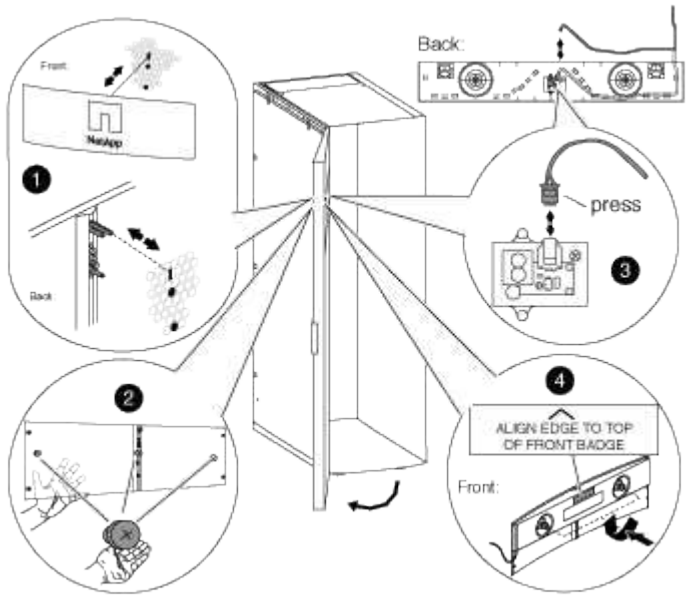

2. Close and lock the front door.

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

https://www.netapp.com/company/legal/copyright/

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

https://www.netapp.com/company/legal/trademarks/

Patents

A current list of NetApp owned patents can be found at:

https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf

Privacy policy

https://www.netapp.com/company/legal/privacy-policy/

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

Notice for disk shelves

Safety information and regulatory notices

https://library.netapp.com/ecm/ecm_download_file/ECMP12475945

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.